arXiv:1407.0312v3 [cs.IT] 19 Nov 2014

Identifying Outliers in Large Matrices via Randomized Adaptive Compressive Sampling

Xingguo Li and Jarvis Haupt

***Abstract*—This paper examines the problem of locating outlier columns in a large, otherwise low-rank, matrix. We propose a simple two-step adaptive sensing and inference approach and establish theoretical guarantees for its performance; our results show that accurate outlier identification is achievable using very few linear summaries of the original data matrix – as few as the squared rank of the low-rank component plus the number of outliers, times constant and logarithmic factors. We demonstrate the performance of our approach experimentally in two stylized applications, one motivated by robust collaborative filtering tasks, and the other by saliency map estimation tasks arising in computer vision and automated surveillance, and also investigate extensions to settings where the data are noisy, or possibly incomplete.**

***Index Terms*—Adaptive sensing, compressed sensing, robust PCA, sparse inference**

I. Introduction

In this paper we address a matrix *outlier identification* problem. Suppose $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2}$ is a data matrix that admits a decomposition of the form

$$\mathbf{M} = \mathbf{L} + \mathbf{C},$$

where $\mathbf{L}$ is a low-rank matrix, and $\mathbf{C}$ is a matrix of outliers that is nonzero in only a fraction of its columns. We are ultimately interested in identifying the locations of the nonzero columns of $\mathbf{C}$, with a particular focus on settings where $\mathbf{M}$ may be very large. The question we address here is, can we accurately (and efficiently) identify the locations of the outliers from a small number of linear measurements of $\mathbf{M}$?

Our investigation is motivated in part by robust collaborative filtering applications, in which the goal may be to identify the locations (or even quantify the number) of corrupted data points or outliers in a large data array. Such tasks may arise in a number of contemporary applications, for example, when identifying malicious responses in survey data or anomalous patterns in network traffic, to name a few. Depending on the nature of the outliers, conventional low-rank approximation approaches based on principal component analysis (PCA) [1], [2] may be viable options for these tasks, but such approaches become increasingly computationally demanding as the data become very high-dimensional. Here, our aim is to leverage dimensionality reduction ideas along the lines of those utilized in randomized numerical linear algebra, (see, e.g., [3], [4] and the references therein) and compressed sensing (see, e.g., [5]–[7]), in order to reduce the size of the data on which our approach operates. In so doing, we also reduce the computational burden of the inference approach relative to comparable methods that operate on "full data."

We are also motivated by an image processing task that arises in many computer vision and surveillance applications – that of identifying the "saliency map" [8] of a given image, which (ideally) indicates the regions of the image that tend to attract the attention of a human viewer. Saliency map estimation is a well-studied area, and numerous methods have been proposed for obtaining saliency maps for a given image – see, for example, [9]–[13]. In contrast to these (and other) methods designed to identify saliency map of an image as a "post processing" step, our aim here is to estimate the saliency map *directly from compressive samples* – i.e., without first performing full image reconstruction as an intermediate step. We address this problem here using a linear subspace-based model of saliency, wherein we interpret an image as a collection of distinct (non-overlapping) patches, so that images may be (equivalently) represented as matrices whose columns are *vectorized* versions of the patches. Previous efforts have demonstrated that such local patches extracted from natural images may be well *approximated* as vectors in a union of low-dimensional linear subspaces (see, e.g., [14]). Here, our approach to the saliency map estimation problem is based on an assumption that salient regions in an image may be modeled as outliers from a single common low-dimensional subspace; the efficacy of similar saliency models for visual saliency has been established recently in [15]. Our approach here may find utility in rapid threat detection in security and surveillance applications in high-dimensional imaging tasks where the goal is not to image the entire scene, but rather to merely identify regions in the image space corresponding to anomalous behavior. Successful identification of salient regions could comprise a first step in an active vision task, where subsequent imaging is restricted to the identified regions.

A. Innovations and Our Approach

We propose a framework that employs dimensionality reduction techniques within the context of a two-step adaptive sampling and inference procedure, and our approach is based on a few key insights. First, we exploit the fact that the enabling geometry of our problem (to be formalized in the following section) is approximately preserved if we operate not on $\mathbf{M}$ directly, but instead on a "compressed" version $\mathbf{\Phi M}$ that has potentially many fewer rows. Next, we use the fact that we can learn the (ostensibly, low-dimensional) linear subspace spanned by the columns of the low rank component of $\mathbf{\Phi M}$

Submitted June 30, 2014; revised October 12, 2014. The authors are with the Department of Electrical and Computer Engineering at the University of Minnesota – Twin Cities. Tel/fax: (612) 625-3300 / (612) 625-4583. Emails: {lixx1661, jdhaupt}@umn.edu. (Corresponding author: J. Haupt.) A version of this work was submitted to ICASSP 2015. The authors graciously acknowledge support from the NSF under Award No. CCF-1217751.

Algorithm 1 Adaptive Compressive Outlier Sensing (ACOS)

Assume: $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2}$

Input: Column sampling Bernoulli parameter $\gamma \in [0,1]$, regularization parameter $\lambda > 0$, Measurement matrices $\mathbf{\Phi} \in \mathbb{R}^{m \times n_1}$, $\mathbf{A} \in \mathbb{R}^{p \times n_2}$, measurement vector $\boldsymbol{\phi} \in \mathbb{R}^{1 \times m}$

Initalize: Column sampling matrix $\mathbf{S} = \mathbf{I}_{:,\mathcal{S}}$, where $\mathcal{S} = \{i : S_i = 1\}$ with $\{S_i\}_{i \in [n_2]}$ i.i.d. Bernoulli(γ)

Step 1

Collect Measurements: $\mathbf{Y}_{(1)} = \mathbf{\Phi M S}$

Solve: $\{\widehat{\mathbf{L}}_{(1)}, \widehat{\mathbf{C}}_{(1)}\} = \underset{\boldsymbol{L},\boldsymbol{C}}{\operatorname{argmin}} \|\boldsymbol{L}\|_* + \lambda \|\boldsymbol{C}\|_{1,2}$ s.t. $\mathbf{Y}_{(1)} = \boldsymbol{L} + \boldsymbol{C}$

Let: $\widehat{\mathcal{L}}_{(1)}$ be the linear subspace spanned by col's of $\widehat{\mathbf{L}}_{(1)}$

Step 2

Compute: $\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}}$, the orthogonal projector onto $\widehat{\mathcal{L}}_{(1)}$

Set: $\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^{\perp}} \triangleq \mathbf{I} - \mathbf{P}_{\widehat{\mathcal{L}}_{(1)}}$

Collect Measurements: $\mathbf{y}_{(2)} = \boldsymbol{\phi}\, \mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^{\perp}} \mathbf{\Phi M A}^T$

Solve: $\widehat{\mathbf{c}} = \underset{\mathbf{c}}{\operatorname{argmin}} \ \|\mathbf{c}\|_1 \ \text{ s.t. } \mathbf{y}_{(2)} = \mathbf{c}\mathbf{A}^T$

Output: $\widehat{\mathcal{I}}_{\mathbf{C}} = \{i : \widehat{\mathrm{c}}_i \neq 0\}$

Algorithm 2 Simplified ACOS (SACOS)

Assume: $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2}$

Input: Column sampling Bernoulli parameter $\gamma \in [0,1]$, regularization parameter $\lambda > 0$, Measurement matrices $\mathbf{\Phi} \in \mathbb{R}^{m \times n_1}$, $\mathbf{A} \in \mathbb{R}^{p \times n_2}$, measurement vector $\boldsymbol{\phi} \in \mathbb{R}^{1 \times m}$

Initalize: Column sampling matrix $\mathbf{S} = \mathbf{I}_{:,\mathcal{S}}$, where $\mathcal{S} = \{i : S_i = 1\}$ with $\{S_i\}_{i \in [n_2]}$ i.i.d. Bernoulli(γ)

Step 1

Collect Measurements: $\mathbf{Y} = \mathbf{\Phi M}$

Form: $\mathbf{Y}_{(1)} = \mathbf{Y S}$

Solve: $\{\widehat{\mathbf{L}}_{(1)}, \widehat{\mathbf{C}}_{(1)}\} = \underset{\boldsymbol{L},\boldsymbol{C}}{\operatorname{argmin}} \|\boldsymbol{L}\|_* + \lambda \|\boldsymbol{C}\|_{1,2}$ s.t. $\mathbf{Y}_{(1)} = \boldsymbol{L} + \boldsymbol{C}$

Let: $\widehat{\mathcal{L}}_{(1)}$ be the linear subspace spanned by col's of $\widehat{\mathbf{L}}_{(1)}$

Step 2

Compute: $\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}}$, the orthogonal projector onto $\widehat{\mathcal{L}}_{(1)}$

Set: $\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^{\perp}} \triangleq \mathbf{I} - \mathbf{P}_{\widehat{\mathcal{L}}_{(1)}}$

Form: $\mathbf{Y}_{(2)} = \mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^{\perp}} \mathbf{Y}$

Form: $\widehat{\mathbf{c}}$ with $\widehat{c}_i = \|(\mathbf{Y}_{(2)})_{:,i}\|_2$ for all $i \in [n_2]$

Output: $\widehat{\mathcal{I}}_{\mathbf{C}} = \{i : \widehat{\mathrm{c}}_i \neq 0\}$

using a small, randomly selected subset of the columns of $\mathbf{\Phi M}$. Our algorithmic approach for this step utilizes a recently proposed method called *Outlier Pursuit* (OP) [16] that aims to separate a matrix $\mathbf{Y}$ into its low-rank and column-sparse components using the convex optimization

$$\underset{\boldsymbol{L},\boldsymbol{C}}{\operatorname{argmin}} \quad \|\boldsymbol{L}\|_* + \lambda \|\boldsymbol{C}\|_{1,2} \ \text{ s.t. } \mathbf{Y} = \boldsymbol{L} + \boldsymbol{C} \tag{1}$$

where $\|\boldsymbol{L}\|_*$ denotes the nuclear norm of $\boldsymbol{L}$ (the sum of its singular values), $\|\boldsymbol{C}\|_{1,2}$ is the sum of the ℓ_2 norms of the columns of $\boldsymbol{C}$, and $\lambda > 0$ is a regularization parameter. Finally, we leverage the fact that correct identification of the subspace spanned by the low-rank component of $\mathbf{\Phi M}$ facilitates (simple) inference of the column outliers.

We analyze two variants of this overall approach. The first (depicted as Algorithm 1) is based on the notion that, contingent on correct identification of the subspace spanned by the low-rank component of $\mathbf{\Phi M}$, we may effectively transform the overall outlier identification problem into a compressed sensing problem, using a carefully-designed linear measurement operator whose net effect is to *(i)* reduce the overall $n_1 \times n_2$ matrix to a $1 \times n_2$ vector whose elements are (nominally) nonzero only at the locations of the outlier columns, and *(ii)* compressively sample the resulting vector. This reduction enables us to employ well-known theoretical results (e.g., [17]) to facilitate our overall analysis. We call this approach Adaptive Compressive Outlier Sensing (ACOS).

The second approach, which we call Simplified ACOS (SACOS) and summarize as Algorithm 2, foregoes the additional dimensionality reduction in the second step and identifies as outliers those columns of $\mathbf{\Phi M}$ having a nonzero component orthogonal to the subspace spanned by the low-rank component of $\mathbf{\Phi M}$. The simplified approach has a (perhaps significantly) higher sample complexity than ACOS, but (as we will see in Section IV) benefits from an ability to identify a larger number of outlier columns relative to the ACOS method. In effect, this provides a trade-off between detection performance and sample complexity for the two methods.

B. Related Work

Our effort here leverages results from Compressive Sensing (CS), where parsimony in the object or signal being acquired, in the form of *sparsity*, is exploited to devise efficient procedures for acquiring and reconstructing high-dimensional objects [5]–[7], [17]. The sequential and adaptive nature of our proposed approach is inspired by numerous recent works in the burgeoning area of adaptive sensing and adaptive CS (see, for example, [18]–[35] as well as the summary article [36] and the references therein). The column subsampling inherent in the first step of our approaches is also reminiscent of the data partitioning strategy of the *divide-and-conquer* parallelization approach of [37] (though our approach only utilizes one small partition of the data for the first inference step).

Our efforts here utilize a generalization of the notion of sparsity, formalized in terms of a low-rank plus outlier matrix model. In this sense, our efforts here are related to earlier work in Robust PCA [38], [39] that seek to identify low-rank matrices in the presence of sparse impulsive outliers, and their extensions to settings where the outliers present as entire columns of an otherwise low-rank matrix [16], [40]–[43]. In fact, the computational approach and theoretical analysis of the first step of our approach make direct utilization of the results of [16].

We also note a related work [44], which seeks to decompose matrices exhibiting some simple structure (e.g., low-rank plus sparse, etc.) into their constituent components from compressive observations. Our work differs from that approach in both the measurement model and scope. Namely, our linear measurements take the form of row and column operations on the matrix and our overall approach is adaptive in nature, in contrast to the non-adaptive "global" compressive measure-

ments acquired in [44], each of which is essentially a linear combination of all of the matrix entries. Further, the goal of [44] was to exactly recover the constituent components, while our aim is only to identify the locations of the outliers. We discuss some further connections with [44] in Section VI.

A component of our numerical evaluation here entails assessing the performance of our approach in a stylized image processing task of saliency map estimation. We note that several recent works have utilized techniques from the sparse representation literature in salient region identification, and in compressive imaging scenarios. A seminal effort in this direction was [45], which proposed a model for feature identification via the human visual cortex based on parsimonious (sparse) representations. More recently, [46] applied techniques from *dictionary learning* [45], [47] and low-rank-plus-sparse matrix decomposition [38], [39] in a procedure to identify salient regions of an image from (uncompressed) measurements. Similar sparse representation techniques for salient feature identification were also examined in [48]. An adaptive compressive imaging procedure driven by a saliency "map" obtained via low-resolution discrete cosine transform (DCT) measurements was demonstrated in [49]. Here, unlike in [46], [48], we consider salient feature identification based on compressive samples, and while our approach is similar in spirit to the problem examined in [49], here we provide theoretical guarantees for the performance of our approach. Finally, we note several recent works [50], [51] that propose methods for identifying salient elements in a data set using compressive samples.

C. Outline

The remainder of the paper is organized as follows. In Section II we formalize our problem, state relevant assumptions, and state our main theoretical results that establish performance guarantees for the Adaptive Compressive Outlier Sensing (ACOS) approach of Algorithm 1 and the Simplified ACOS approach of Algorithm 2. In Section III we outline the proofs of our main results. Section IV contains the results of a comprehensive experimental evaluation of our approach on synthetic data, as well as in a stylized image processing application of saliency map estimation. In section V we empirically investigate several extensions of our methods to noisy and "missing data" scenarios. In Section VI we provide a brief discussion of the computational complexity of our approach, and discuss a few potential future directions. We relegate proofs and other auxiliary material to the appendix.

D. A Note on Notation

We use bold-face upper-case letters ($\mathbf{M}, \mathbf{L}, \mathbf{C}, \mathbf{\Phi}, \boldsymbol{L}, \boldsymbol{C}, \boldsymbol{I}$ etc.) to denote matrices, and use the MATLAB-inspired notation $\mathbf{I}_{:,\mathcal{S}}$ to denote the sub matrix formed by extracting columns of $\mathbf{I}$ indexed by $i \in \mathcal{S}$. We typically use bold-face lower-case letters ($\mathbf{x}, \mathbf{v}, \boldsymbol{c}, \boldsymbol{\phi}$, etc.) to denote vectors, with an exception along the lines of the indexing notation above – i.e., that $\mathbf{C}_{:,i}$ denotes the i-th column of $\mathbf{C}$. Note that we employ both "block" and "math" type notation (e.g., $\mathbf{L}$, $\boldsymbol{L}$), where the latter are used to denote variables in the optimization tasks that arise throughout our exposition. Non-bold letters are used to denote scalar parameters or constants; the usage will be made explicit, or will be clear from context.

The ℓ_1 norm of a vector $\mathbf{x} = [\mathrm{x}_1 \ \mathrm{x}_2 \ \ldots \ \mathrm{x}_n]$ is $\|\mathbf{x}\|_1 = \sum_i |\mathrm{x}_i|$ and the ℓ_2 norm is $\|\mathbf{x}\|_2 = \left(\sum_i |\mathrm{x}_i|^2\right)^{1/2}$. We denote the nuclear norm (the sum of singular values) of a matrix $\boldsymbol{L}$ by $\|\boldsymbol{L}\|_*$ and the $1,2$ norm (the sum of column ℓ_2 norms) of a matrix $\boldsymbol{C}$ by $\|\boldsymbol{C}\|_{1,2}$. We denote the operator norm (the largest singular value) of a matrix $\boldsymbol{L}$ by $\|\boldsymbol{L}\|$. Superscript asterisks denote complex conjugate transpose.

For positive integers n, we let $[n]$ denote the set of positive integers no greater than n; that is, $[n] = \{1, 2, \ldots, n\}$.

II. Main Results

A. Problem Statement

Our specific problem of interest here may be formalized as follows. We suppose $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2}$ admits a decomposition of the form $\mathbf{M} = \mathbf{L} + \mathbf{C}$, where $\mathbf{L}$ is a low-rank matrix having rank at most r, and $\mathbf{C}$ is a matrix having some $k \leq n_2$ nonzero columns that we will interpret as "outliers" from $\mathbf{L}$, in the sense that they do not lie (entirely) within the span of the columns of $\mathbf{L}$. Formally, let $\mathcal{L}$ denote the linear subspace of $\mathbb{R}^{n_1}$ spanned by the columns of $\mathbf{L}$ (and having dimension at most r), denote its orthogonal complement in $\mathbb{R}^{n_1}$ by $\mathcal{L}^\perp$, and let $\mathbf{P}_{\mathcal{L}}$ and $\mathbf{P}_{\mathcal{L}^\perp}$ denote the orthogonal projection operators onto $\mathcal{L}$ and $\mathcal{L}^\perp$, respectively. We assume that the nonzero columns of $\mathbf{C}$ are indexed by a set $\mathcal{I}_{\mathbf{C}}$ of cardinality k, and that $i \in \mathcal{I}_{\mathbf{C}}$ if and only if $\|\mathbf{P}_{\mathcal{L}^\perp}\mathbf{C}_{:,i}\|_2 > 0$. Aside from this assumption, the elements of the nonzero columns of $\mathbf{C}$ may be *arbitrary*.

Notice that *without loss of generality*, we may assume that the columns of $\mathbf{L}$ are zero at the locations corresponding to the nonzero columns of $\mathbf{C}$ (since those columns of $\mathbf{L}$ can essentially be aggregated into the nonzero columns of $\mathbf{C}$, and the resulting column will still be an outlier according to our criteria above). We adopt that model here, and assume $\mathbf{L}$ has a total of $n_{\mathbf{L}}$ nonzero columns[1], including all k of the indices in $\mathcal{I}_{\mathbf{C}}$ where $\mathbf{C}$ has a nonzero column, but also potentially others, to allow for the case where some $n_{\mathbf{L}} - k$ columns of $\mathbf{M}$ itself to be zero. Clearly $n_{\mathbf{L}} \leq n_2 - k$.

Given this setup, our problem of interest here may be stated concisely – our aim is to identify the set $\mathcal{I}_{\mathbf{C}}$ containing the locations of the outlier columns.

B. Assumptions

It is well-known in the matrix completion and robust PCA literature that separation of low-rank and sparse matrices from observations of their sum may not be a well-posed task – for example, matrices having only a single nonzero element are simultaneously low rank, sparse, column-sparse, row-sparse, etc. To overcome these types of identifiability issues, it is common to assume that the linear subspace spanned by the

[1]As we will see, the conditions under which our column subsampling in Step 1 succeeds will depend on the number of *nonzero* columns in the low-rank component, since any all-zero columns are essentially non-informative for learning the low-rank subspace. Thus, we make the distinction between n_2 and $n_{\mathbf{L}}$ explicit throughout.

rows and/or columns of the low-rank matrix be "incoherent" with the canonical basis (see, e.g., [16], [38]–[40], [52]).

In a similar vein, since our aim is to identify column outliers from an otherwise low-rank matrix we seek conditions that make the factors distinguishable so that any of the directions of the column space of $\mathbf{L}$ that we seek to identify are not defined by a single vector (stated another way, we would like the vectors whose columns comprise $\mathbf{L}$ to be "spread out" in the subspace spanned by columns of $\mathbf{L}$). To this end, we assume an incoherence condition on the row space of the low-rank component $\mathbf{L}$. We formalize this notion via the following definition from [16].

Definition II.1 (Column Incoherence Property)**.** *Let $\mathbf{L} \in \mathbb{R}^{n_1 \times n_2}$ be a rank r matrix with at most $n_{\mathbf{L}} \leq n_2$ nonzero columns, and compact singular value decomposition (SVD) $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$, where $\mathbf{U}$ is $n_1 \times r$, $\mathbf{\Sigma}$ is $r \times r$, and $\mathbf{V}$ is $n_2 \times r$. The matrix $\mathbf{L}$ is said to satisfy the* column incoherence property *with parameter $\mu_{\mathbf{L}}$ if*

$$\max_i \|\mathbf{V}^* \mathbf{e}_i\|_2^2 \leq \mu_{\mathbf{L}} \frac{r}{n_{\mathbf{L}}},$$

where $\{\mathbf{e}_i\}$ are basis vectors of the canonical basis for $\mathbb{R}^{n_2}$.

Note that $\mu_{\mathbf{L}} \in [1, n_{\mathbf{L}}/r]$; the lower limit is achieved when all elements of $\mathbf{V}^*$ have the same amplitude, and the upper limit when any one element of $\mathbf{V}^*$ is equal to 1 (i.e., when the row space of $\mathbf{L}$ is aligned with the canonical basis). For our purposes, an undesirable case occurs when $\mathbf{V}^*$ is such that $\max_i \|\mathbf{V}^* \mathbf{e}_i\|_2^2 = 1$, since this implies that (at least) one of the directions in the span of the columns of $\mathbf{L}$ is described by only a single vector (and thus distinguishing that vector from a column outlier becomes ambiguous).

With this, we may state our assumptions concisely, as follows: we assume that the components $\mathbf{L}$ and $\mathbf{C}$ of the matrix $\mathbf{M} = \mathbf{L} + \mathbf{C}$ satisfy the following *structural conditions*:

- **(c1)** $\operatorname{rank}(\mathbf{L}) = r$,
- **(c2)** $\mathbf{L}$ has $n_{\mathbf{L}}$ nonzero columns,
- **(c3)** $\mathbf{L}$ satisfies the *column incoherence property* with parameter $\mu_{\mathbf{L}}$, and
- **(c4)** $|\mathcal{I}_{\mathbf{C}}| = k$, where $\mathcal{I}_{\mathbf{C}} = \{i : \|\mathbf{P}_{\mathcal{L}^\perp} \mathbf{C}_{:,i}\|_2 > 0, \mathbf{L}_{:,i} = \mathbf{0}\}$.

C. Recovery Guarantees and Implications

Our main results identify conditions under which the procedures outlined in Algorithm 1 and Algorithm 2 succeed. Our particular focus is on the case where the measurement matrices are random, and satisfy the following property.

Definition II.2 (Distributional Johnson-Lindenstrauss (JL) Property)**.** *An $m \times n$ matrix $\mathbf{\Phi}$ is said to satisfy the* distributional JL property *if for any fixed $\mathbf{v} \in \mathbb{R}^n$ and any $\epsilon \in (0, 1)$,*

$$\Pr\left(\left| \|\mathbf{\Phi}\mathbf{v}\|_2^2 - \|\mathbf{v}\|_2^2 \right| \geq \epsilon \|\mathbf{v}\|_2^2 \right) \leq 2e^{-mf(\epsilon)}, \qquad (2)$$

where $f(\epsilon) > 0$ is a constant depending only on ϵ that is specific to the distribution of $\mathbf{\Phi}$.

Random matrices satisfying the distributional JL property are those that preserve the length of any fixed vector to within a multiplicative factor of $(1 \pm \epsilon)$ with probability at least $1 - 2e^{-mf(\epsilon)}$. By a simple union bounding argument, such matrices can be shown to approximately preserve the lengths of a finite collection of vectors, all vectors in a linear subspace, all vectors in a union of subspaces, etc., provided the number of rows is sufficiently large. As noted in [53], for many randomly constructed and appropriately normalized $\mathbf{\Phi}$, (e.g., such that entries of $\mathbf{\Phi}$ are i.i.d. zero-mean Gaussian, or are drawn as an ensemble from any subgaussian distribution), $f(\epsilon)$ is quadratic[2] in ϵ as $\epsilon \to 0$. This general framework also allows us to directly utilize other specially constructed *fast* or *sparse* JL transforms [55], [56].

With this, we are in position to formulate our first main result. We state it here as a theorem; its proof appears in Section III.

Theorem II.1 (Accurate Recovery via ACOS)**.** *Suppose $\mathbf{M} = \mathbf{L} + \mathbf{C}$, where the components $\mathbf{L}$ and $\mathbf{C}$ satisfy the structural conditions **(c1)**-**(c4)** with*

$$k \leq \frac{1}{40(1 + 121\ r\mu_{\mathbf{L}})}\ n_2. \qquad (3)$$

For any $\delta \in (0, 1)$, if the column subsampling parameter γ satisfies

$$\gamma \geq \max\left\{ \frac{1}{20}, \frac{200 \log(\frac{5}{\delta})}{n_{\mathbf{L}}}, \frac{24 \log(\frac{10}{\delta})}{n_2}, \frac{10 r \mu_{\mathbf{L}} \log(\frac{5r}{\delta})}{n_{\mathbf{L}}} \right\}, \qquad (4)$$

the measurement matrices are each drawn from any distribution satisfying (2) with

$$m \geq \frac{5(r+1) + \log(k) + \log(2/\delta)}{f(1/4)} \qquad (5)$$

and

$$p \geq \frac{11k + 2k \log(n_2/k) + \log(2/\delta)}{f(1/4)}, \qquad (6)$$

the elements of $\boldsymbol{\phi}$ are i.i.d. realizations of any continuous random variable, and for any upper bound k_{ub} of k the regularization parameter is set to $\lambda = \frac{3}{7\sqrt{k_{\text{ub}}}}$, then the following hold simultaneously with probability at least $1 - 3\delta$:

- *the ACOS procedure in Algorithm 1 correctly identifies the salient columns of $\boldsymbol{C}$ (i.e., $\widehat{\mathcal{I}}_{\boldsymbol{C}} = \mathcal{I}_{\boldsymbol{C}}$), and*
- *the total number of measurements collected is no greater than $\left(\frac{3}{2}\right) \gamma m n_2 + p$.*

It is interesting to compare this result with that of [16], which established that the Outlier Pursuit procedure (1) succeeds in recovering the true low-rank subspace and locations of the outlier columns provided $\mathbf{M}$ satisfy conditions analogous to **(c1)**-**(c4)** with $k \leq n_2/(1 + (121/9)\ r\mu_{\mathbf{L}})$. The sufficient condition (3) on the number of recoverable outliers that we identify for the ACOS procedure differs from the condition identified in that work by only constant factors. Further, the number of identifiable outliers could be as large as a fixed fraction of n_2 when both the rank r and coherence parameter $\mu_{\mathbf{L}}$ are small.

It is also interesting to note the sample complexity improvements that are achievable using the ACOS procedure.

[2]It was shown in [54], for example, that $f(\epsilon) = \epsilon^2/4 - \epsilon^3/6$ for matrices whose elements are appropriately normalized Gaussian or symmetric Bernoulli random variables.

Namely, it follows directly from our analysis that for appropriate choice of the parameters γ, m, and p the ACOS algorithm correctly identifies the salient columns of $\mathbf{C}$ with high probability from relatively few observations, comprising only a fraction of the measurements required by other comparable (non-compressive) procedures [16] that produce the same correct salient support estimate but operate directly on the full $(n_1 \times n_2)$ matrix $\boldsymbol{M}$. Specifically, our analysis shows that the ACOS approach succeeds with high probability with an effective sampling rate of $\frac{\#\text{obs}}{n_1 n_2} = \mathcal{O}\left(\max\left\{ \frac{(r+\log k)(n_2/n_{\mathbf{L}})\mu_{\mathbf{L}} r \log r}{n_1 n_2}, \frac{(r+\log k)}{n_1} \right\} + \frac{k \log(n_2/k)}{n_1 n_2} \right)$, which may be small when r and k are each small relative to the problem dimensions (and $n_{\mathbf{L}} \sim n_2$, so that $\mathbf{L}$ does not have a large number of zero columns outside of $\mathcal{I}_{\mathbf{C}}$).

Another point of comparison for our result comes from the related work [40], which addresses a different (and in a sense, more difficult) task of identifying both the column space and the set of outlier columns of a matrix $\mathbf{M} = \mathbf{L} + \mathbf{C}$ from observations that take the form of samples of the elements of $\mathbf{M}$. There, to deal with the fact that observations take the form of point samples of the matrix (rather than more general linear measurements as here), the authors of [40] assume that $\mathbf{L}$ also satisfy a row incoherence property in addition to a column incoherence property, and show that in this setting that the column space of $\mathbf{L}$ and set of nonzero columns of $\mathbf{C}$ may be recovered from only $\mathcal{O}\left(n_2 r^2 \mu^2 \log(n_2)\right)$ observations via a convex optimization, where $\mu \in [1, n_1/r]$ is the row incoherence parameter. Normalizing this sample complexity by $n_1 n_2$ facilitates comparison with our result above; we see that the sufficient conditions for the sample complexity of our approach are smaller than for the approach of [40] by a factor of at least $1/r$, and, our approach does not require the row incoherence assumption. We provide some additional, experimental, comparisons between our ACOS method and the RMC method in Section IV.

We may also obtain performance guarantees for Algorithm 2 (in effect, using a simplified version of the analysis used to establish Theorem II.1). This yields the following corollary.

Corollary II.1 (Accurate Recovery via SACOS)**.** *Suppose $\mathbf{M} = \mathbf{L} + \mathbf{C}$, where the components $\mathbf{L}$ and $\mathbf{C}$ satisfy the structural conditions **(c1)**-**(c4)** with k as in (3). Let the measurement matrix $\boldsymbol{\Phi}$ be drawn from a distribution satisfying (2), and assume (4) and (5) hold. If for any upper bound k_{ub} of k the regularization parameter is set to $\lambda = \frac{3}{7\sqrt{k_{\text{ub}}}}$, then the following hold simultaneously with probability at least $1 - 2\delta$:*

- *the ACOS procedure in Algorithm 2 correctly identifies the salient columns of $\boldsymbol{C}$ (i.e., $\widehat{\mathcal{I}}_{\boldsymbol{C}} = \mathcal{I}_{\boldsymbol{C}}$), and*
- *the total number of measurements collected is no greater than mn_2.*

We leave the proof (which is straightforward, using the lemmata in the following section) to the interested reader.

III. Proof of Theorem II.1

First, we note that in both of the steps of Algorithm 1 the prescribed observations are functions of $\mathbf{M}$ only through $\boldsymbol{\Phi}\mathbf{M}$; stated another way, $\mathbf{M}$ never appears in the algorithm in isolation from the measurement matrix $\boldsymbol{\Phi}$. Motivated by this, we introduce

$$\widetilde{\mathbf{M}} \triangleq \boldsymbol{\Phi}\mathbf{M} = \boldsymbol{\Phi}\mathbf{L} + \boldsymbol{\Phi}\mathbf{C} = \widetilde{\mathbf{L}} + \widetilde{\mathbf{C}}, \tag{7}$$

to effectively subsume the action of $\boldsymbol{\Phi}$ into $\widetilde{\mathbf{M}}$. Now, our proof is a straightforward consequence of assembling three intermediate probabilistic results via a union bounding argument. The first intermediate result establishes that for $\mathbf{M} = \mathbf{L} + \mathbf{C}$ with components $\mathbf{L}$ and $\mathbf{C}$ satisfying the structural conditions **(c1)**-**(c4)**, the components $\widetilde{\mathbf{L}}$ and $\widetilde{\mathbf{C}}$ of $\widetilde{\mathbf{M}}$ as defined in (7) satisfy analogous structural conditions provided that m, the number of rows of $\boldsymbol{\Phi}$, be sufficiently large. We state this result here as a lemma; its proof appears in Appendix A.

Lemma III.1. *Suppose $\mathbf{M} = \mathbf{L} + \mathbf{C}$, where $\mathbf{L}$ and $\mathbf{C}$ satisfy the structural conditions **(c1)**-**(c4)**. Fix any $\delta \in (0,1)$, suppose $\boldsymbol{\Phi}$ is an $m \times n_1$ matrix drawn from a distribution satisfying the distributional JL property (2) with m satisfying (5) and let $\widetilde{\mathbf{M}} = \widetilde{\mathbf{L}} + \widetilde{\mathbf{C}}$ be as defined in (7). Then, the components $\widetilde{\mathbf{L}}$ and $\widetilde{\mathbf{C}}$ satisfy the following conditions simultaneously with probability at least $1 - \delta$:*

- **($\widetilde{\boldsymbol{c}}$1)** $\operatorname{rank}(\widetilde{\mathbf{L}}) = r$,
- **($\widetilde{\boldsymbol{c}}$2)** *$\widetilde{\mathbf{L}}$ has $n_{\mathbf{L}}$ nonzero columns,*
- **($\widetilde{\boldsymbol{c}}$3)** *$\widetilde{\mathbf{L}}$ satisfies the column incoherence property with parameter $\mu_{\mathbf{L}}$, and*
- **($\widetilde{\boldsymbol{c}}$4)** *$\mathcal{I}_{\widetilde{\mathbf{C}}} \triangleq \{i : \|\mathbf{P}_{\widetilde{\mathcal{L}}^\perp} \widetilde{\mathbf{C}}_{:,i}\|_2 > 0, \widetilde{\mathbf{L}}_{:,i} = \mathbf{0}\} = \mathcal{I}_{\mathbf{C}}$, where $\widetilde{\mathcal{L}}$ is the linear subspace of $\mathbb{R}^m$ spanned by the columns of $\widetilde{\mathbf{L}}$, and $\mathbf{P}_{\widetilde{\mathcal{L}}^\perp}$ denotes the orthogonal projection onto the orthogonal complement of $\widetilde{\mathcal{L}}$ in $\mathbb{R}^m$.*

The second intermediate result guarantees two outcomes – first, that Step 1 of Algorithm 1 succeeds in identifying the correct column space of $\widetilde{\mathcal{L}}$ (i.e., that $\widehat{\mathcal{L}}_{(1)} = \widetilde{\mathcal{L}}$) with high probability provided the components $\widetilde{\mathbf{L}}$ and $\widetilde{\mathbf{C}}$ of $\widetilde{\mathbf{M}}$ as specified in (7) satisfy the structural conditions **($\widetilde{\boldsymbol{c}}$1)**-**($\widetilde{\boldsymbol{c}}$4)** and the column sampling probability parameter γ be sufficiently large, and second, that the number of columns of the randomly generated sampling matrix $\mathbf{S}$ be close to γn_2. We also provide this result as a lemma; its proof appears in Appendix B.

Lemma III.2. *Let $\widetilde{\mathbf{M}} = \widetilde{\mathbf{L}} + \widetilde{\mathbf{C}}$ be an $m \times n_2$ matrix, where the components $\widetilde{\mathbf{L}}$ and $\widetilde{\mathbf{C}}$ satisfy the conditions **($\widetilde{\boldsymbol{c}}$1)**-**($\widetilde{\boldsymbol{c}}$4)** with k satisfying (3). Fix $\delta \in (0,1)$ and suppose the column sampling parameter γ satisfies (4). When $\lambda = \frac{3}{7\sqrt{k_{\text{ub}}}}$ for any $k_{\text{ub}} \geq |\mathcal{I}_{\widetilde{\mathbf{C}}}|$, the following hold simultaneously with probability at least $1 - \delta$: $\mathbf{S}$ has $|\mathcal{S}| \leq (3/2)\gamma n_2$ columns, and the subspace $\widehat{\mathcal{L}}_{(1)}$ resulting from Step 1 of Algorithm 1 satisfies $\widehat{\mathcal{L}}_{(1)} = \widetilde{\mathcal{L}}$.*

Our third intermediate result shows that the support set of the vector $\widehat{\mathbf{c}}$ produced in Step 2 of Algorithm 1 is the same as the set of salient columns of $\widetilde{\mathbf{C}}$, provided that $\widehat{\mathcal{L}}_{(1)} = \widetilde{\mathcal{L}}$ and that p, the number of rows of $\mathbf{A}$, is sufficiently large. We state this result here as a lemma; its proof appears in Appendix C

Lemma III.3. *$\widetilde{\mathbf{M}} = \widetilde{\mathbf{L}} + \widetilde{\mathbf{C}}$ be an $m \times n_2$ matrix, where the components $\widetilde{\mathbf{L}}$ and $\widetilde{\mathbf{C}}$ satisfy the conditions **($\widetilde{\boldsymbol{c}}$1)**-**($\widetilde{\boldsymbol{c}}$4)** for any $k \leq n_2$, and suppose $\widehat{\mathcal{L}}_{(1)} = \widetilde{\mathcal{L}}$, the subspace spanned by the columns of $\widetilde{\mathbf{L}}$. Let $\boldsymbol{\Phi}\mathbf{M} = \widetilde{\mathbf{M}}$ in Step 2 of Algorithm 1. Fix $\delta \in (0,1)$, suppose $\mathbf{A}$ is a $p \times n_2$ matrix drawn from a distribution satisfying the distributional JL property (2) with p satisfying*

(6), *and suppose the elements of* ϕ *are i.i.d. realizations of any continuous random variable. Then with probability at least* $1-\delta$ *the support* $\mathcal{I}_{\widehat{\mathbf{c}}} \triangleq \{i : \widehat{\mathrm{c}}_i \neq 0\}$ *of the vector* $\widehat{\mathbf{c}}$ *produced by Step 2 of Algorithm 1 satisfies* $\mathcal{I}_{\widehat{\mathbf{c}}} = \mathcal{I}_{\widetilde{\mathbf{C}}}$.

Our overall result follows from assembling these intermediate results via union bound. In the event that the conclusion of Lemma III.1 holds, then so do the requisite conditions of Lemma III.2. Thus, with probability at least $1-2\delta$ the conclusions of Lemmata III.1 and III.2 both hold. This implies that the requisite conditions of Lemma III.3 hold also with probability at least $1-2\delta$, and so it follows that the conclusions of all three Lemmata hold with probability at least $1-3\delta$.

IV. Experimental Evaluation

In this section we provide a comprehensive experimental evaluation of the performance of our approaches for both synthetically generated and real data, the latter motivated by a stylized application of saliency map estimation in an image processing task. We compare our methods with the Outlier Pursuit (OP) approach of [16] and the Robust Matrix Completion (RMC) approach of [40], each of which employs a convex optimization to identify both the subspace in which the columns of the low rank matrix lie, and the locations of the nonzero columns in the outlier matrix. We implement the RMC method using an accelerated approximate alternating direction method of multipliers (ADMM) method inspired by [57] (as well as [16], [58]). We implement the OP methods (as well as the intermediate execution of the OP-like optimization in Step 1 of our approach) using the procedure in [40]. We implement the ℓ_1-regularized estimation in Step 2 of our procedure by casting it as a LASSO problem and using an accelerated proximal gradient method [58].

A. Synthetic Data

We experiment on synthetically generated $n_1 \times n_2$ matrices $\mathbf{M}$, with $n_1 = 100$ and $n_2 = 1000$, formed as follows. For a specified rank r and number of outliers k, we let the number of nonzero columns of $\mathbf{L}$ be $n_{\mathbf{L}} = n_2 - k$, generate two random matrices $\mathbf{U} \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{V} \in \mathbb{R}^{n_{\mathbf{L}} \times r}$ with i.i.d. $\mathcal{N}(0,1)$ entries, and we take $\mathbf{L} = [\mathbf{U}\mathbf{V}^T \ \mathbf{0}_{n_1 \times k}]$. We generate the outlier matrix $\mathbf{C}$ as $\mathbf{C} = [\mathbf{0}_{n_1 \times n_{\mathbf{L}}} \ \mathbf{W}]$ where $\mathbf{W} \in \mathbb{R}^{n_1 \times k}$ has i.i.d. $\mathcal{N}(0,r)$ entries (which are also independent of entries of $\mathbf{U}$ and $\mathbf{V}$). Then, we set $\mathbf{M} = \mathbf{L} + \mathbf{C}$. Notice that the outlier vector elements have been scaled, so that all columns of $\mathbf{M}$ have the same squared ℓ_2 norm, in expectation. In all experiments we generate ϕ, $\mathbf{\Phi}$, and $\mathbf{A}$ with i.i.d. zero-mean Gaussian entries.

Our first experiment investigates the "phase transition" behavior of our ACOS approach; our experimental setting is as follows. First, we set the average sampling rate by fixing the column downsampling fraction $\gamma = 0.2$, and choosing a row sampling parameter $m \in \{0.1n_1, 0.2n_1, 0.3n_1\}$ and column sampling parameter $p \in \{0.1n_2, 0.2n_2, 0.3n_2\}$. Then, for each (r,k) pair with $r \in \{1,2,3,\ldots,40\}$ and $k \in \{2,4,6,\ldots,100\}$ we generate a synthetic matrix $\mathbf{M}$ as above, and for each of 3 different values of the regularization parameter $\lambda \in \{0.3, 0.4, 0.5\}$ we perform 100 trials of Algorithm 1 recording in each whether the recovery approach succeeded[3] in identifying the locations of the true outliers for that value of λ, and associate to each (r,k) pair the (empirical) average success rate. Then, at each (r,k) point examined we identify the point-wise maximum of the average success rates for the 3 different values of λ; in this way, we assess whether recovery for that (r,k) is achievable by our method for the specified sampling regime for *some* choice of regularization parameters. The results in Figure 1 depict the outcome of this experiment for the 9 different sampling regimes examined. For easy comparison, we provide the average sampling rate as fraction of observations obtained (relative to the full matrix dimension) in the caption in each figure.

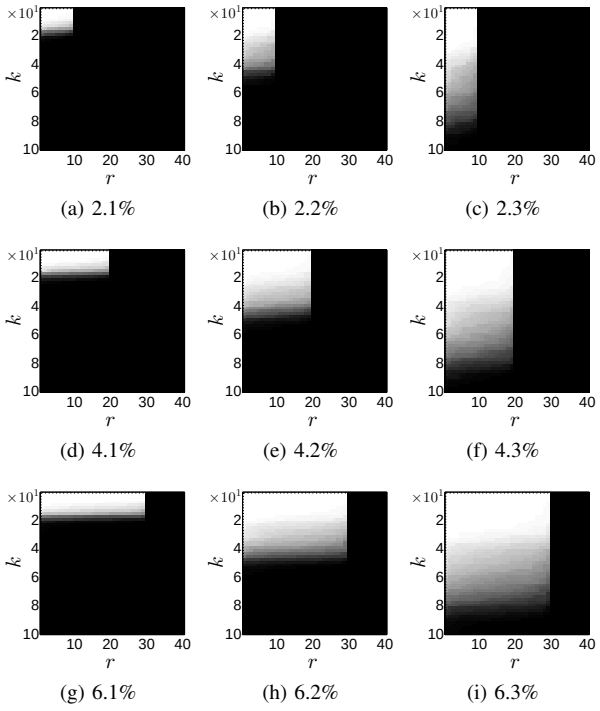

Fig. 1. Outlier recovery phase transitions plots for ACOS (white regions correspond to successful recovery). Each row of the figure corresponds to a different level of compression of rows of $\mathbf{M}$, where $m = 0.1n_1, 0.2n_1$ and $0.3n_1$, respectively, from top to bottom. Each column corresponds to a different level of compression of rows of $\mathbf{M}$ in Step 2 of Algorithm 1, with $p = 0.1n_2, 0.2n_2$ and $0.3n_2$, respectively, from left to right. The fraction of observations obtained (as a percentage, relative to the full dimension) is provided as a caption below each figure. As expected, increasing m (top to bottom) facilitates accurate estimation for increasing rank r of $\mathbf{L}$, while increasing p (left to right) allows for recovery of increasing numbers k of outlier columns.

The results of this experiment provide an interesting, and somewhat intuitive, illustration of the efficacy of our approach. Namely, we see that increasing the parameter m of the matrix $\mathbf{\Phi}$ in Step 1 of our algorithm while keeping the other sampling parameters fixed (i.e., moving from top to bottom in any one column) facilitates accurate recovery for increasing ranks r

[3] We solve the optimization associated with Step 2 of our approach as a LASSO problem, with 10 different choices of regularization parameter $\mu \in (0,1)$. We deem any trial a success if for at least one value of μ, there exists a threshold $\tau > 0$ such that $\min_{i \in \mathcal{I}_{\mathbf{C}}} |\widehat{\mathrm{c}}_i(\mu)| > \tau > \max_{j \notin \mathcal{I}_{\mathbf{C}}} |\widehat{\mathrm{c}}_j(\mu)|$ for the estimate $\widehat{\mathbf{c}}(\mu)$ produced in Step 2. An analogous threshold-based methodology was employed to assess the outlier detection performance of the Outlier Pursuit approach in [16].

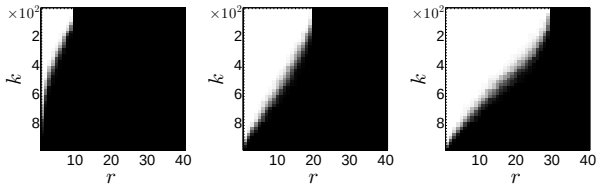

(a) 10% (b) 20% (c) 30%

Fig. 2. Outlier recovery phase transitions plots for SACOS (white regions correspond to successful recovery). The row sampling parameters are $m = 0.1n_1$, $0.2n_1$, and $0.3n_1$ respectively, from left to right. Increasing m in SACOS enables accurate estimation for larger rank and increasing numbers of outlier columns. The sampling rate is provided below each plot.

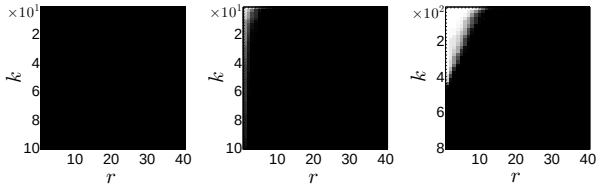

(a) 5% (RMC) (b) 10% (RMC) (c) 20% (RMC)

Fig. 3. Outlier recovery phase transitions plots for RMC. The average sampling rates are 5%, 10% and 20%, from left to right. Note that the vertical (k) scale in panels (a) and (b) matches that of Figure 1, while the scale on panel (c) matches that of Figure 2. Further, comparing panels (a) and (b) here with Figure 1 shows that ACOS outperforms RMC at low sampling rates, while comparing panels (b) and (c) here with panels (a) and (b) of Figure 2 shows that SACOS yields correct outlier identification for a larger portion of the parameter space than RMC for the same average sampling rates.

of the matrix $\mathbf{L}$. Similarly, increasing the parameter p of the matrix $\mathbf{A}$ in Step 2 of our algorithm while keeping the other sampling parameters fixed (i.e., moving from left to right in any one row) facilitates accurate recovery for an increasing number k of outlier columns. Overall, our approach can successfully recover the locations of the outliers for non-trivial regimes of r and k using very few measurements – see, for instance, panel (i), where $\sim$ 30 outlier columns can be accurately identified in the presence of a rank ~ 30 background using an effective sampling rate of only $\sim 6.3\%$.

We adopt a similar methodology to evaluate the Simplified ACOS approach, except that we set $k \in \{20, 40, 60, \ldots, 980\}$ (and the parameter p is no longer applicable, since there is no additional compression in Step 2 for this method). The results are shown in Figure 2. As noted above the SACOS approach has a higher average sampling rate than ACOS for the same m, but the results show this facilitates recovery of much larger numbers k of outlier columns (notice the difference in the vertical scales in Figures 1 and 2). Overall, we may view ACOS and SACOS as complementary; when the number k of outlier columns is relatively small and low sampling ratio $\frac{\#\text{obs}}{n_1 n_2}$ is a primary focus, ACOS may be preferred, while if the number k of outlier columns is relatively large, SACOS is more favorable (at the cost of increased sample complexity).

We also compute phase transition curves for RMC using a similar methodology to that described above. The results are provided in Figure 3 . We observe[4] that RMC approach is viable for identifying the outliers from subsampled data provided the sampling rate exceeds about 10%, but even then only for small values of the rank r. As alluded in the discussion in previous sections, the relative difference in performance is likely due in large part to the difference in the observation models between the two approaches – the RMC approach is inherently operating in the presence of "missing data" (a difficult scenario!) while our approach permits us to observe linear combinations of any row or column of the entire matrix (i.e., we are allowed to "see" each entry of the matrix, albeit not necessarily individually, throughout our approach).

[4]Our evaluation of RMC here agrees qualitatively with results in [40], where sampling rates around 10% yielded successful recovery for small r.

B. Real Data

We also evaluate the performance of our proposed methods on real data in the context of a stylized image processing task that arises in many computer vision and automated surveillance – that of identifying the "saliency map" of an image. For this, we use images from the MSRA Salient Object Database [12], available online at http://research.microsoft.com/en-us/um/people/jiansun/SalientObject/salient_object.htm.

As discussed above, our approach here is based on representing each test image as a collection of (vectorized) non-overlapping image patches. We transform each (color) test image to gray scale, decompose it into non-overlapping 10×10-pixel patches, vectorize each patch into a 100×1 column vector, and assemble the column vectors into a matrix. Most of the images in the database are of the size 300×400 (or 400×300), which here yields matrices of size 100×1200, corresponding to 1200 patches. Notice that we only used gray scale values of image as the input feature rather than any high-level images feature – this facilitates the use of our approach, which is based on collecting linear measurements of the data (e.g., using a spatial light modulator, or an architecture like the *single pixel camera* [59]).

Here, our experimental approach is (somewhat necessarily) a bit more heuristic than for the synthetic data experiments above, due in large part to the fact that the data here may not adhere exactly to the low-rank plus outlier model. To compensate for this, we augment Step 1 of Algorithm 1 and Algorithm 2 with an additional "rank reduction" step, where we further reduce the dimension of the subspace spanned by the columns of the learned $\widehat{\mathbf{L}}_{(1)}$ by truncating its SVDs to retain the smallest number of leading singular values whose sum is at least $0.95 \times \|\widehat{\mathbf{L}}_{(1)}\|_*$. Further, we generalize Step 2 of each procedure by declaring an image patch to be salient when its (residual) column norm is sufficiently large, rather than strictly nonzero. We used visual heuristics to determine the "best" outputs for Step 2 of each method, selecting LASSO parameters (for ACOS) or thresholds (for SACOS) in order to qualitatively trade off false positives with misses.

We implement our ACOS and SACOS methods using three different sampling regimes for each, with the fixed column downsampling parameter $\gamma = 0.2$ throughout. For ACOS, we examine settings where $m = 0.2n_1$, $0.1n_1$ and $0.05n_1$ with $p = 0.5n_2$, which result in average sampling rates of 4.5%, 2.5% and 1.5%, respectively. For SACOS, we examine settings where $m = 0.2n_1$, $0.05n_1$ and $0.03n_1$, resulting in average sampling rates of 20%, 5% and 3%, respectively.

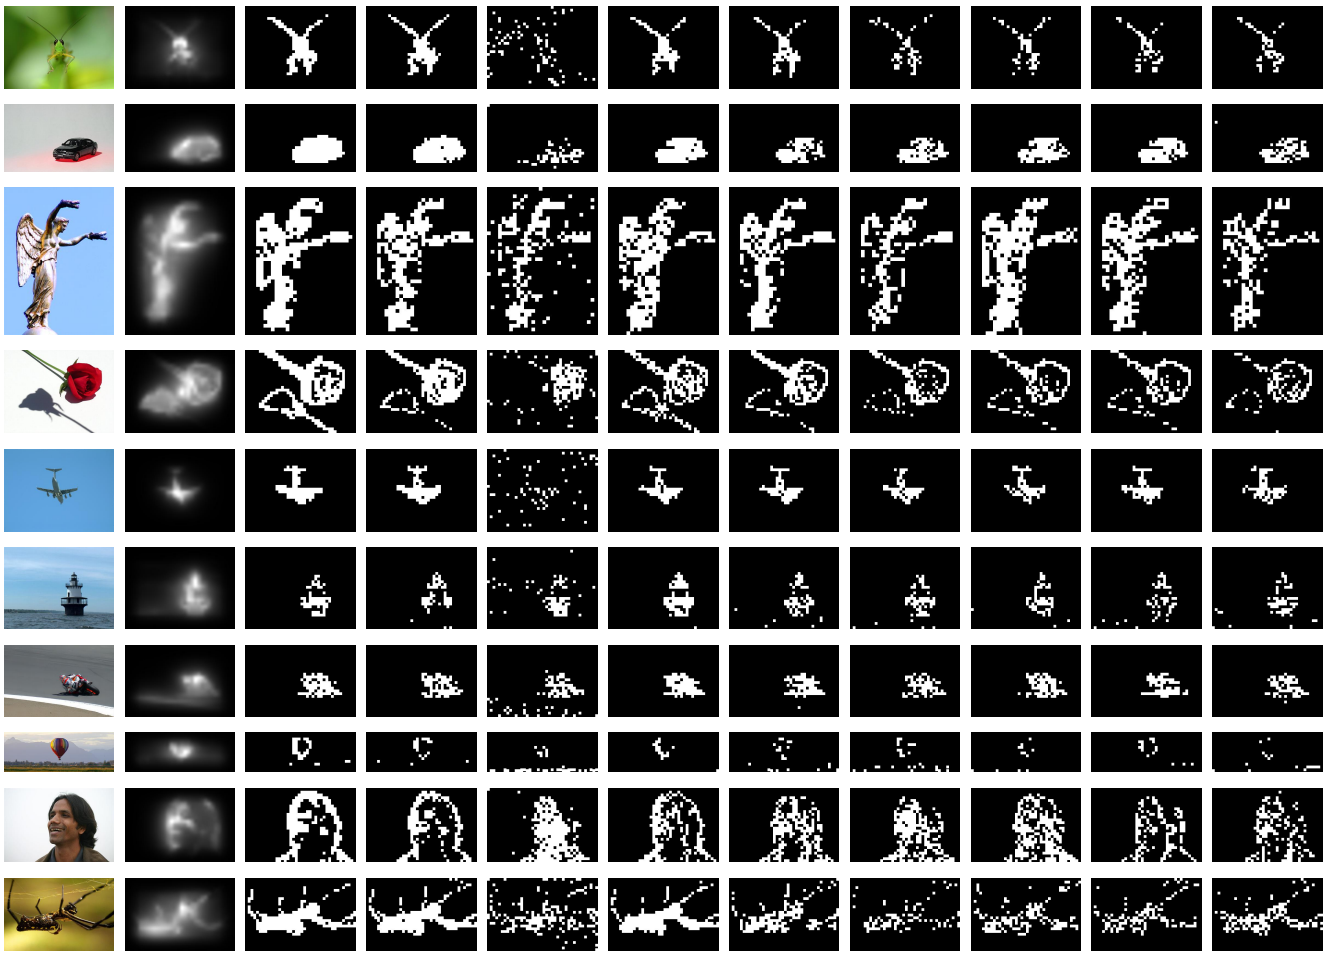

Fig. 4. Detection results for the MSRA Salient Object Database for various methods. Our ACOS approach produces results comparable to the "full sampling" OP method using an average sampling rate below 5%. The performance of the RMC approach appears to degrade at low sampling rates.

TABLE I
TIMING ANALYSIS FOR DETECTION EXPERIMENTS ON 1000 IMAGES FROM MRSA DATABASE.
EACH ENTRY IS THE MEAN EXECUTION TIME IN SECONDS WITH THE STANDARD DEVIATION IN PARENTHESIS.

Method Sampling	GBVS 100%	OP 100%	RMC 20%	RMC 5%	SACOS 20%	SACOS 5%	SACOS 3%	ACOS 4.5%	ACOS 2.5%	ACOS 1.5%
Step 1	0.9926 (0.2742)	2.9441 (0.3854)	2.6324 (0.3237)	2.7254 (0.3660)	0.0538 (0.0121)	0.0107 (0.0034)	0.0074 (0.0017)	0.0533 (0.0118)	0.0214 (0.0056)	0.0105 (0.0025)
Step 2	– –	– –	– –	– –	0.0015 (0.0003)	0.0011 (0.0003)	0.0009 (0.0003)	0.2010 (0.0674)	0.2014 (0.0692)	0.2065 (0.0689)

As before, we generate the $\mathbf{\Phi}$ and $\mathbf{A}$ matrices to have i.i.d. zero-mean Gaussian entries. We compare our approaches with two "benchmarks" – the Graph-based visual saliency (GBVS) method from the computer vision literature [11] and the OP approach (both of which use the full data) – as well as with the RMC approach at sampling rates of 20% and 5%.

The results of this experiment are provided in Figure 4. We note first that the OP approach performs fairly well at identifying the visually salient regions in the image, essentially identifying the same salient regions as the GBVS procedure and providing evidence to validate the use of the low-rank plus outlier model for visual saliency (see also [15]). Next, comparing the results of the individual procedures, we see that the OP approach appears to uniformly give the best detection results, which is reasonable since it is using the full data as input. The RMC approach performs well at the 20% sampling rate, but its performance appears to degrade at the 5% sampling rate. The SACOS approach, on the other hand, still produces reasonably accurate results using only 3% sampling. Moreover, ACOS provides acceptable results even with 1-2% sampling rate.

We also compare implementation times of the algorithms on this saliency map estimation task. Table I provides the average execution times (and standard deviations) for each approach, evaluated over 1000 images in the MSRA database[5]. Here, we only execute each procedure for one choice of regularization parameter, and we also include the additional

[5]Timing comparisons were done with MATLAB R2013a on an iMac with a 3.4 GHz Intel Core i7 processor, 32 GB memory, and running OS X 10.8.5.

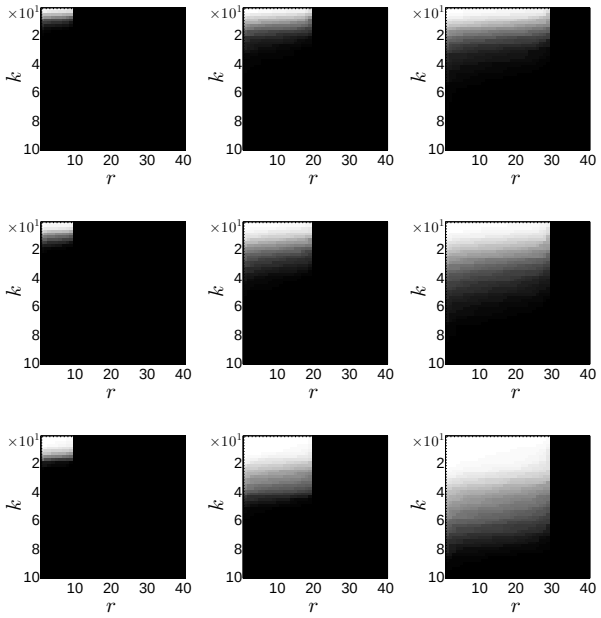

(a) 2.1% (b) 4.2% (c) 6.3%

Fig. 5. Outlier recovery phase transitions plots for ACOS for noisy settings (white regions correspond to successful recovery). Rows correspond to $\sigma = 0.001$, 0.0005 and 0.0001 respectively, from top to bottom; columns correspond to the settings $m = 0.1n_1,\ p = 0.1n_2$; $m = 0.2n_1,\ p = 0.2n_2$; and $m = 0.3n_1,\ p = 0.3n_2$ respectively, from left to right. The fraction of observations obtained is provided below each figure. As in Figure 1, larger m and p promote accurate recovery for increasing rank r and numbers k of outlier columns. Here, however, increasing noise variance degrades the estimation results, especially with respect to the number k of outliers that can be accurately identified.

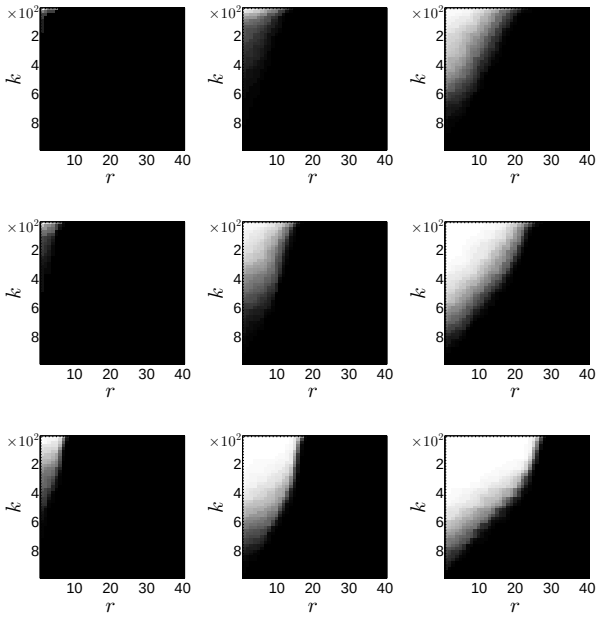

(a) 10% (b) 20% (c) 30%

Fig. 6. Outlier recovery phase transitions plots for SACOS for noisy settings (white regions correspond to successful recovery). Rows of the figure correspond to $\sigma = 0.03$, 0.02 and 0.01 respectively, from top to bottom; columns correspond to $m = 0.1n_1,\ 0.2n_1$, and $0.3n_1$ respectively, from left to right. The fraction of observations obtained is provided below each column. In this case, increasing noise variance results in a decrease in both the rank r as well as the number k of outliers that can be accurately identified.

"rank reduction" step discussed above for the ACOS and SACOS methods. Overall, we see the ACOS approach is up to $4\times$ faster than the GBVS method and $15\times$ faster than the OP and RMC methods, while the SACOS approach could result overall in relative speedups of $100\times$ over GBVS and $300\times$ over the OP and RMC methods. Overall, our results suggest a significant improvement obtained via ACOS and SACOS for both detection consistency and timing, which may have a promising impact in a variety of salient signal detection tasks.

V. Extensions

A. Noisy Observations

We demonstrate the outlier detection performance of our approaches under the scenario when $\mathbf{M}$ is contaminated by unknown random noise or modeling error. Formally, we consider the setting where $\mathbf{L}$ and $\mathbf{C}$ are as above, but

$$\mathbf{M} = \mathbf{L} + \mathbf{C} + \mathbf{N}, \tag{8}$$

where $\mathbf{N}$ has i.i.d. $\mathcal{N}(0, \sigma^2)$ entries.

We first investigate the performance of the ACOS method, following a similar experimental methodology as in Section IV to generate $\mathbf{L}$ and $\mathbf{C}$, except that now we renormalize each column of $(\mathbf{L}+\mathbf{C})$ to have unit Euclidean norm (essentially to standardize the noise levels). We consider three different noise levels ($\sigma = 0.001$, 0.0005 and 0.0001), three pairs of the row sampling parameter m and the column sampling parameter p ($m = 0.1n_1,\ p = 0.1n_2$; $m = 0.2n_1,\ p = 0.2n_2$; and $m = 0.3n_1,\ p = 0.3n_2$) and for each we fix the column downsampling fraction to be $\gamma = 0.2$; the corresponding sampling ratios are 2.1%, 4.2% and 6.3%, respectively. We again perform 100 trials of Algorithm 1 and record the success frequency for each. The results are given in Figure 5.

It can be observed from the results that increasing m and p promote accurate estimation of outlier column indices for increasing rank r and numbers k of outlier columns, which is exactly what we have seen in Figure 1 for the noiseless case. However, the presence of noise degrades the estimation performance, albeit gracefully. This is reasonable, since in Step 2 of Algorithm 1, the measurements $\mathbf{y}_2$ might be perturbed more seriously as the energy of noise increases, which results in more difficult recovery of true supports of $\mathbf{c}$. Under this scenario, we will require larger p to enable better recovery of the underlying true supports.

We also evaluate the SACOS procedure in noisy settings for three choices of m ($m = 0.1n_1$, $0.2n_1$ and $0.3n_2$) and fixed column downsampling fraction $\gamma = 0.2$. Here, we again normalize columns of $(\mathbf{L}+\mathbf{C})$, but consider three higher noise levels, corresponding to $\sigma = 0.03$, 0.02 and 0.01. The results are presented in Figure 6. Here, we again observe a graceful performance degradation with noise. Notice, however, that higher level of variances of noise can be tolerated for SACOS compared with ACOS, which is an artifact of the difference between the second (inference) steps of the two procedures.

B. Missing Data

We also describe and demonstrate an extension of our SACOS method that is amenable to scenarios characterized by missing data. Suppose that there exists some underlying matrix $\mathbf{M}$ that admits a decomposition of the form $\mathbf{M} = \mathbf{L} + \mathbf{C}$ with $\mathbf{L}$ and $\mathbf{C}$ as above, but we are only able to observe $\mathbf{M}$ at a subset of its locations. Formally, we denote by $\mathbf{\Omega} \subseteq [n_1] \times [n_2]$ the set of indices corresponding to the available elements of $\mathbf{M}$, and let $\mathbf{P}_{\mathbf{\Omega}}(\cdot)$ be the operator that masks its argument at locations not in $\mathbf{\Omega}$. Thus, rather than operate on $\mathbf{M}$ itself, we consider procedures that operate on the sampled data $\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M})$.

In this setting, we can modify our SACOS approach so that the observations obtained in Step 1 are of the form $\mathbf{Y}_{(1)} = \mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M})\mathbf{S}$, where (as before) $\mathbf{S}$ is a column selection matrix but $\mathbf{\Phi}$ is now a *row* subsampling matrix (i.e., it is comprised of a subset of rows of the $n_1 \times n_1$ identity matrix) containing some m rows. The key insight here is that the composite operation of sampling elements of $\mathbf{M}$ followed by row subsampling can be expressed in terms of a related operation of subsampling elements of a row-subsampled version of $\mathbf{M}$. Specifically, we have that $\mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M}) = \mathbf{P}_{\mathbf{\Omega}_{\mathbf{\Phi}}}(\mathbf{\Phi}\mathbf{M})$, where $\mathbf{P}_{\mathbf{\Omega}_{\mathbf{\Phi}}}(\cdot)$ masks the same elements as $\mathbf{P}_{\mathbf{\Omega}}(\cdot)$ in the rows selected by $\mathbf{\Phi}$.

Now, given $\mathbf{Y}_{(1)}$, we solve a variant of RMC [40]

$$\{\widehat{\mathbf{L}}_{(1)}, \widehat{\mathbf{C}}_{(1)}\} = \underset{\boldsymbol{L},\boldsymbol{C}}{\operatorname{argmin}} \|\boldsymbol{L}\|_* + \lambda \|\boldsymbol{C}\|_{1,2} \quad \text{s.t. } \mathbf{Y}_{(1)} = \mathbf{P}_{\mathbf{\Omega}_{\mathbf{\Phi}}}(\boldsymbol{L} + \boldsymbol{C})$$

in an initial step, identifying (as before) an estimate $\widehat{\mathbf{L}}_{(1)}$ whose column span is an estimate of the subspace spanned by the low-rank component of $\mathbf{\Phi}\mathbf{M}$.

Then (in a second step) we perform the "missing data" analog of the orthogonal projection operation on every column $j \in [n_2]$ of $\mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M})$, as follows. For each $j \in [n_2]$, we let $\mathcal{I}_j \in [m]$ denote the locations at which observations of column j of $\mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M})$ are available, and let $(\mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M}))_{\mathcal{I}_j,j}$ be the sub vector of $(\mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M}))_{:,j}$ containing only the elements indexed by $\mathcal{I}_j$. Similarly, let $(\widehat{\mathbf{L}}_{(1)})_{\mathcal{I}_j,:}$ be the row submatrix of $\widehat{\mathbf{L}}_{(1)}$ formed by retaining rows indexed by $\mathcal{I}_j$. Now, let $\mathbf{P}_{\widehat{\mathcal{L}}_{(1)_j}}$ denote the orthogonal projection onto the subspace spanned by columns of $(\widehat{\mathbf{L}}_{(1)})_{\mathcal{I}_j,:}$ and compute the residual energy of the j-th column as $\|(\mathbf{I} - \mathbf{P}_{\widehat{\mathcal{L}}_{(1)_j}})(\mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M}))_{\mathcal{I}_j,j}\|_2$. Overall, the orthogonal projection for the j-th column of $\mathbf{\Phi}\mathbf{P}_{\mathbf{\Omega}}(\mathbf{M})$ is only computed over the nonzero entries of that column, an approach motivated by a recent effort examining subsampling methods in the context of matrix completion [60].

We evaluate this approach empirically using the same data generation methods as above, and using an independent *Bernoulli* model to describe the subsampling operation $\mathbf{P}_{\mathbf{\Omega}}(\cdot)$ (so that each $(i,j) \in \mathbf{\Omega}$ independently with probability $p_{\mathbf{\Omega}}$). We consider noise-free settings, fix the column subsampling parameter $\gamma = 0.2$, and examine three different row-sampling scenarios ($m = 0.1n_1$, $0.2n_1$ and $0.3n_1$) in each choosing subsets of m rows uniformly at random from the collection of all $\binom{n_1}{m}$ sets of cardinality m. The results are in Figure 7. Again, increasing m and p permits accurate estimation of outlier column indices for increasing rank r and numbers k of outlier columns. Further, we do observe the performance degradation as the number of missing entries of $\mathbf{M}$ increases.

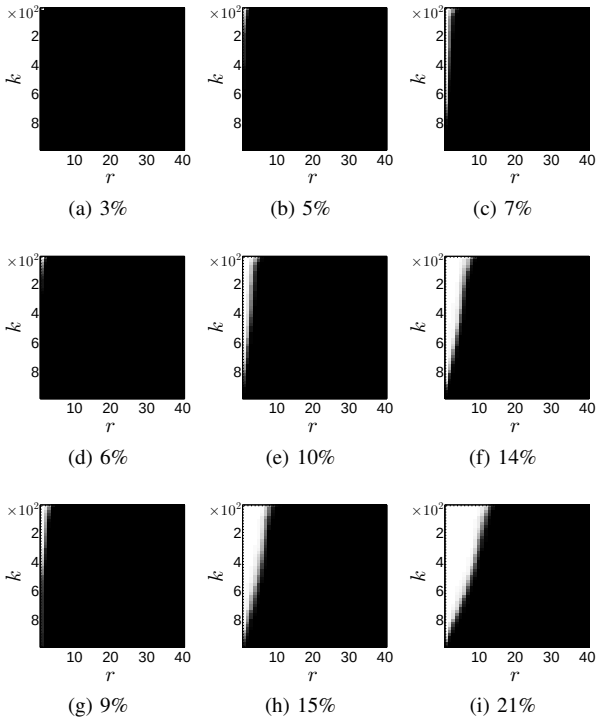

Fig. 7. Outlier recovery phase transitions plots for a "missing data" variant of the SACOS method (white regions correspond to successful recovery). Rows correspond to available data fractions of $p_{\mathbf{\Omega}}$=0.3, 0.5 and 0.7 respectively, from top to bottom; columns corresponds to row sampling parameters $m = 0.1n_1$, $0.2n_1$, and $0.3n_1$, respectively, from left to right.

VI. Discussion and Future Directions

It is illustrative here to note a key difference between our approach and more conventional compressive sensing (CS) tasks. Namely, the goal of the original CS works [5]–[7] and numerous follow-on efforts was to exactly recover or reconstruct a signal from compressive measurements, whereas the nature of our task here is somewhat simpler, amounting to a kind of multidimensional "support recovery" task (albeit in the presence of a low-rank "background"). Exactly recovering the low-rank and column-sparse components would be sufficient for the outlier identification task we consider here, but as our analysis shows it is not strictly necessary. This is the insight that we exploit when operating on the "compressed" data $\mathbf{\Phi}\mathbf{M}$ instead of the original data matrix $\mathbf{M}$. Ultimately, this allows us to successfully identify the locations of the outliers *without first estimating the original (full size) low-rank matrix or the outliers themselves*. For some regimes of $\mu_{\mathbf{L}}$, r and k, we accomplish the outlier identification task using as few as $\mathcal{O}\left((r + \log k)(\mu_{\mathbf{L}} r \log r) + k \log(n_2/k)\right)$ observations.

Along related lines, it is reasonable to conjecture that any procedure would require at least $r^2 + k$ measurements in order to identify k outliers from an r-dimensional linear subspace. Indeed, a necessary condition for the existence of outliers of a rank-r subspace, as we have defined them, is that the number of rows of $\mathbf{M}$ be at least $r+1$. Absent any additional structural conditions on the outliers and the subspace spanned

TABLE II
COMPUTATIONAL COMPLEXITIES OF OUTLIER IDENTIFICATION METHODS. THE STATED RESULTS ASSUME USE OF AN ACCELERATED FIRST ORDER METHOD FOR ALL SOLVERS (SEE TEXT FOR ADDITIONAL DETAILS).

Method	Complexity
OP	$\mathcal{O}\left(\mathrm{IT}\cdot[n_1n_2\cdot\min\{n_1,n_2\}]\right)$
RMC	$\mathcal{O}\left(\mathrm{IT}\cdot[n_1n_2\cdot\min\{n_1,n_2\}]\right)$
ACOS	$\mathcal{O}\left(\mathrm{IT}_1\,[m(\gamma n_2)\min\{m,\gamma n_2\}]+\mathrm{IT}_2\,[pn_2]\right)$
SACOS	$\mathcal{O}\left(\mathrm{IT}_1\,[m(\gamma n_2)\min\{m,\gamma n_2\}]+m^2n_2\right)$

by columns of the low-rank matrix, one would need to identify a collection of r vectors that span the r-dimensional subspace containing the column vectors of the low-rank component (requiring specification of some $\mathcal{O}(r^2)$ parameters) as well as the locations of the k outliers (which would entail specifying another k parameters). In this sense, our approach may be operating near the sample complexity limit for this problem, at least for some regimes of $\mu_{\mathbf{L}}$, r and k.

It would be interesting to see whether the dimensionality reduction insight that we exploit in our approach could be leveraged in the context of the Compressive Principal Component Pursuit (Compressive PCP) of [44] in order to yield a procedure with comparable performance as ours, but which acquires only *non-adaptive* linear measurements of $\mathbf{M}$. Direct implementation of that approach in our experimental setting was somewhat computationally prohibitive (e.g., simulations at a 10% sampling rate would require generation and storage of random matrices having 10^9 elements). Alternatively, it is interesting to consider implementing the Compressive PCP method not on the full data $\mathbf{M}$, but on the a priori compressed data $\mathbf{\Phi M}$. Our Lemma III.1 establishes that the row compression step preserves rank and column incoherence properties, so it is plausible that the Compressive PCP approach may succeed in recovering the components of the compressed matrix, which would suffice for the outlier identification task. We defer this investigation along these lines to a future effort.

We also comment briefly on the computational complexities of the methods we examined. We consider first the OP and RMC approaches, and assume that the solvers for each utilize an iterative accelerated first-order method (like those mentioned in the first part of Section IV). In this case, the computational complexity will be dominated by SVD steps in each iteration. Now, for an $n_1\times n_2$ matrix the computational complexity of the SVD is $\mathcal{O}(n_1n_2\cdot\min\{n_1,n_2\})$; with this, and assuming some IT iterations are used, we have that the complexities of both OP and RMC scale as $\mathcal{O}\left(\mathrm{IT}\cdot[n_1n_2\cdot\min\{n_1,n_2\}]\right)$. By a similar analysis, we can conclude that the complexity of Step 1 of the ACOS and SACOS methods scales like $\mathcal{O}\left(\mathrm{IT}_1\cdot[m(\gamma n_2)\cdot\min\{m,\gamma n_2\}]\right)$, where IT_1 denotes the number of iterations for the solver in Step 1. If we further assume an iterative accelerated first-order method for the LASSO in Step 2 of the ACOS approach, and that IT_2 iterations are used, then the second step of the ACOS approach would have overall computational complexity $\mathcal{O}\left(\mathrm{IT}_2\cdot[pn_2]\right)$. Along similar lines, Step 2 of SACOS would entail $\mathcal{O}(m^2n_2+mn_2)=\mathcal{O}(m^2n_2)$ operations to compute the orthogonal projections and their ℓ_2 norms. We summarize the overall complexity results in Table II. Since we will typically have γ small, $m\ll n_1$, and $p\ll n_2$ in our approaches, the computational complexity of our approaches can be much less than methods that operate on the full data or require intermediate SVD's of matrices of the same size as $\mathbf{M}$.

Note that we have not included here the complexity of acquiring or forming the observations in any of the methods. For the ACOS method, this would comprise up to an additional $\mathcal{O}\left(mn_1(\gamma n_2)\right)$ operations for Step 1 and $\mathcal{O}(m^2+mn_1+n_1n_2+n_2p)=\mathcal{O}(mn_1+n_1n_2+n_2p)$ operations for Step 2, where the complexity for the second step is achieved by iteratively multiplying together the left-most two factors in the overall product, and using the fact that $m\leq n_1$. Similarly, observations obtained via the SACOS approach could require up $\mathcal{O}(mn_1n_2)$ operations. On the other hand, depending on the implementation platform, forming the observations themselves could also have a negligible computational effect e.g., in our imaging example when linear observations are formed "implicitly" using a spatial light modulator or single pixel camera [59]. Finally, we note that further reductions in the overall complexity of our approach may be achieved using fast or sparse JL embeddings along the lines of [55], [56].

Finally, it is worth noting[6] that the performance in our our visual saliency application could likely be improved using an additional assumption that the salient regions be spatially clustered. This could be implemented here using *group sparse* regularization (e.g. [61]) in Step 2 of ACOS, or (more simply) by directly identifying groups of nonzero elements in Step 2 of SACOS. We defer investigations along these lines to a future effort.

APPENDIX

A. Proof of Lemma III.1

We proceed using the formalism of *stable embeddings* that has emerged from the dimensionality reduction and compressive sensing literature (see, e.g., [62]).

Definition A.1 (Stable Embedding)**.** *For $\epsilon\in[0,1]$ and $\mathcal{U},\mathcal{V}\subseteq\mathbb{R}^n$, we say $\mathbf{\Phi}$ is an* ϵ-stable embedding *of $(\mathcal{U},\mathcal{V})$ if*

$$(1-\epsilon)\|\mathbf{u}-\mathbf{v}\|_2^2\leq\|\mathbf{\Phi u}-\mathbf{\Phi v}\|_2^2\leq(1+\epsilon)\|\mathbf{u}-\mathbf{v}\|_2^2 \quad (9)$$

for all $\mathbf{u}\in\mathcal{U}$ and $\mathbf{v}\in\mathcal{V}$.

Our proof approach is comprised of two parts. First, we show that each of the four claims in the lemma follow when $\mathbf{\Phi}$ is an ϵ-stable embedding of

$$(\mathcal{L},\cup_{i\in\mathcal{I}_{\mathbf{C}}}\{\boldsymbol{C}_{:,i}\}\cup\{\mathbf{0}\}) \quad (10)$$

for any choice of $\epsilon<1/2$. Second, we show that for any $\delta\in(0,1)$, generating $\mathbf{\Phi}$ as a random matrix as specified in the lemma ensures it will be a $\sqrt{2}/4$-stable embedding of (10) with probability at least $1-\delta$. The choice of $\sqrt{2}/4$ in the last step is somewhat arbitrary – we choose this fixed value for concreteness here, but note that the structural conclusions of the lemma follow using any choice of $\epsilon<1/2$ (albeit with slightly different conditions on m).

[6]Thanks to David B. Dunson and Alfred O. Hero for these suggestions.

1) Part 1: Throughout this portion of the proof we assume that $\mathbf{\Phi}$ is an ϵ-stable embedding of (10) for some $\epsilon < 1/2$, and establish each of the four claims in turn. First, to establish that $\text{rank}(\mathbf{\Phi L}) = r = \text{rank}(\mathbf{L})$, we utilize an intermediate result of [53], stated here as a lemma (without proof) and formulated in the language of stable embeddings.

Lemma A.1 (Adapted from [53], Theorem 1)**.** *Let $\mathbf{L}$ be an $n_1 \times n_2$ matrix of rank r, and let $\mathcal{L}$ denote the column space of $\mathbf{L}$, which is an r-dimensional linear subspace of $\mathbb{R}^{n_1}$. If for some $\epsilon \in (0,1)$, $\mathbf{\Phi}$ is an ϵ-stable embedding of $(\mathcal{L}, \{\mathbf{0}\})$ then* $\text{rank}(\mathbf{\Phi L}) = r = \text{rank}(\mathbf{L})$.

Here, since $\mathbf{\Phi}$ being an ϵ-stable embedding of (10) implies it is also an ϵ-stable embedding of $(\mathcal{L}, \{\mathbf{0}\})$, the first claim (of Lemma III.1) follows from Lemma A.1.

Next we show that $\mathbf{\Phi L}$ has $n_{\mathbf{L}}$ nonzero columns. Since $\mathbf{\Phi}$ is a stable embedding of $(\mathcal{L}, \{\mathbf{0}\})$, it follows that for each of the $n_{\mathbf{L}}$ nonzero columns $\mathbf{L}_{:,i}$ of $\mathbf{L}$ we have $\|\mathbf{\Phi L}_{:,i}\|_2^2 > (1-\epsilon)\|\mathbf{L}_{:,i}\|_2^2 > 0$, while for each of the remaining $n_2 - n_{\mathbf{L}}$ columns $\mathbf{L}_{:,j}$ of $\mathbf{L}$ that are identically zero we have $\|\mathbf{\Phi L}_{:,j}\|_2^2 = 0$ so that $\mathbf{\Phi L}_{:,j} = 0$.

Continuing, we show next that $\mathbf{\Phi L}$ satisfies the column incoherence property with parameter $\mu_{\mathbf{L}}$. Recall from above that we write the compact SVD of $\mathbf{L}$ as $\mathbf{L} = \mathbf{U\Sigma V}^*$, where $\mathbf{U}$ is $n_1 \times r$, $\mathbf{V}$ is $n_2 \times r$, and $\mathbf{\Sigma}$ is an $r \times r$ nonnegative diagonal matrix of singular values (all of which are strictly positive). The incoherence condition on $\mathbf{L}$ is stated in terms of column norms of the matrix $\mathbf{V}^*$ whose rows form an orthonormal basis for the row space of $\mathbf{L}$. Now, when the rank of $\mathbf{\Phi L}$ is the same as that of $\mathbf{L}$, which is true here on account of Lemma A.1, the row space of $\mathbf{\Phi L}$ is *identical* to that of $\mathbf{L}$, since each are r-dimensional subspaces of $\mathbb{R}^{n_2}$ spanned by linear combinations of the columns of the $\mathbf{V}^*$. Thus since the rank and number of nonzero columns of $\mathbf{\Phi L}$ are the same as for $\mathbf{L}$, the coherence parameter of $\mathbf{\Phi L}$ is just $\mu_{\mathbf{L}}$, and the third claim is established.

Finally, we establish the last claim, that the set of salient columns of $\mathbf{\Phi C}$ is the same as for $\mathbf{C}$. Recall that the condition that a column $\mathbf{C}_{:,i}$ be salient was equivalent to the condition that $\|\mathbf{P}_{\mathcal{L}^\perp}\mathbf{C}_{:,i}\|_2 > 0$, where $\mathbf{P}_{\mathcal{L}^\perp}$ is the orthogonal projection operator onto the orthogonal complement of $\mathcal{L}$ in $\mathbb{R}^{n_1}$. Here, our aim is to show that an analogous result holds in the "projected" space – that for all $i \in \mathcal{I}_{\mathbf{C}}$ we have $\|\mathbf{P}_{(\mathbf{\Phi}\mathcal{L})^\perp}\mathbf{\Phi C}_{:,i}\|_2 > 0$, where $\mathbf{\Phi}\mathcal{L}$ is the linear subspace spanned by the columns of $\mathbf{\Phi L}$. For this we utilize an intermediate result of [62] formulated there in terms of a "compressive interference cancellation" method. We state an adapted version of that result here as a lemma (without proof).

Lemma A.2 (Adapted from [62], Theorem 5)**.** *Let $\mathcal{V}_1$ be an r-dimensional linear subspace of $\mathbb{R}^n$ with $r < n$, let $\mathcal{V}_2$ be any subset of $\mathbb{R}^n$, and let $\check{\mathcal{V}}_2 = \{\mathbf{P}_{\mathcal{V}_1^\perp}\mathbf{v} : \mathbf{v} \in \mathcal{V}_2\}$, where $\mathbf{P}_{\mathcal{V}_1^\perp}$ is the orthogonal projection operator onto the orthogonal complement of $\mathcal{V}_1$ in $\mathbb{R}^n$. If $\mathbf{\Phi}$ is an ϵ-stable embedding of $(\mathcal{V}_1, \check{\mathcal{V}}_2 \cup \{\mathbf{0}\})$, then for all $\check{\mathbf{v}} \in \check{\mathcal{V}}_2$*

$$\|\mathbf{P}_{(\mathbf{\Phi}\mathcal{V}_1)^\perp}(\mathbf{\Phi}\check{\mathbf{v}})\|_2^2 \geq \left(1 - \frac{\epsilon}{1-\epsilon}\right)\|\check{\mathbf{v}}\|_2^2, \tag{11}$$

where $\mathbf{P}_{(\mathbf{\Phi}\mathcal{V}_1)^\perp}$ is the orthogonal projection operator onto the orthogonal complement of the subspace of $\mathbb{R}^n$ spanned by the elements of $\mathbf{\Phi}\mathcal{V}_1 = \{\mathbf{\Phi v} : \mathbf{v} \in \mathcal{V}_1\}$.

Before applying this result we first note a useful fact, that $\mathbf{\Phi}$ being an ϵ-stable embedding of $(\mathcal{V}_1, \check{\mathcal{V}}_2 \cup \{\mathbf{0}\})$ is equivalent to $\mathbf{\Phi}$ being an ϵ-stable embedding of $(\mathcal{V}_1, \mathcal{V}_2 \cup \{\mathbf{0}\})$, which follows directly from the definition of stable embeddings and the (easy to verify) fact that $\{\mathbf{v}_1 - \check{\mathbf{v}}_2 : \mathbf{v}_1 \in \mathcal{V}_1, \check{\mathbf{v}}_2 \in \check{\mathcal{V}}_2 \cup \{\mathbf{0}\}\} = \{\mathbf{v}_1 - \mathbf{v}_2 : \mathbf{v}_1 \in \mathcal{V}_1, \mathbf{v}_2 \in \mathcal{V}_2 \cup \{\mathbf{0}\}\}$. Now, to apply Lemma A.2 here, we let $\mathcal{V}_1 = \mathcal{L}$, $\mathcal{V}_2 = \cup_{i \in \mathcal{I}_{\mathbf{C}}}\{\mathbf{C}_{:,i}\}$, and $\check{\mathcal{V}}_2 = \cup_{i \in \mathcal{I}_{\mathbf{C}}}\{\mathbf{P}_{\mathcal{L}^\perp}\mathbf{C}_{:,i}\}$. Since $\mathbf{\Phi}$ is an ϵ-stable embedding of (10), we have that for all $i \in \mathcal{I}_{\mathbf{C}_{:,i}}$, $\|\mathbf{P}_{(\mathbf{\Phi}\mathcal{L})^\perp}(\mathbf{\Phi C})_{:,i}\|_2^2 \geq \left(1 - \frac{\epsilon}{1-\epsilon}\right)\|\mathbf{P}_{\mathcal{L}^\perp}\mathbf{C}_{:,i}\|_2^2$. Since $\epsilon < 1/2$, the above result implies $\|\mathbf{P}_{(\mathbf{\Phi}\mathcal{L})^\perp}\mathbf{\Phi C}_{:,i}\|_2 > 0$ for all $i \in \mathcal{I}_{\mathbf{C}}$, while for all $j \notin \mathcal{I}_{\mathbf{C}}$ we have $\mathbf{C}_{:,j} = \mathbf{0}$, implying that $\mathbf{\Phi C}_{:,j} = \mathbf{0}$ and hence $\|\mathbf{P}_{(\mathbf{\Phi}\mathcal{L})^\perp}\mathbf{\Phi C}_{:,j}\|_2 = 0$. Using this, and the fact that the nonzero columns of $\mathbf{\Phi L}$ coincide with the nonzero columns of $\mathbf{L}$, we conclude that $\mathcal{I}_{\mathbf{\Phi C}} = \{i : \|\mathbf{P}_{(\mathbf{\Phi}\mathcal{L})^\perp}\mathbf{\Phi C}_{:,j}\|_2 > 0, (\mathbf{\Phi L})_{:,i} = \mathbf{0}\}$ is the same as $\mathcal{I}_{\mathbf{C}}$.

2) Part 2: Given the structural result established in the previous step, the last part of the proof entails establishing that a random matrix $\mathbf{\Phi}$ generated as specified in the statement of Lemma III.1 is an $\sqrt{2}/4$-stable embedding of (10). Our approach here begins with a brief geometric discussion, and a bit of "stable embedding algebra." Appealing to the definition of stable embeddings, we see that $\mathbf{\Phi}$ being an ϵ-stable embedding of (10) is equivalent to $\mathbf{\Phi}$ being such that

$$(1-\epsilon)\|\mathbf{v}\|_2^2 \leq \|\mathbf{\Phi v}\|_2^2 \leq (1+\epsilon)\|\mathbf{v}\|_2^2 \tag{12}$$

holds for all $\mathbf{v} \in \mathcal{L} \cup \bigcup_{i \in \mathcal{I}_{\mathbf{C}}} \mathcal{L} - \mathbf{C}_{:,i}$, where $\mathcal{L} - \mathbf{C}_{:,i}$ denotes the r-dimensional *affine* subspace of $\mathbb{R}^{n_1}$ comprised of all elements taking the form of a sum between a vector in $\mathcal{L}$ and the fixed vector $\mathbf{C}_{:,i}$. Thus, in words, establishing our claim here entails showing that a random $\mathbf{\Phi}$ (generated as specified in the lemma, with appropriate dimensions) approximately preserves the lengths of all vectors in a *union of subspaces* comprised of one r-dimensional linear subspace and some $|\mathcal{I}_{\mathbf{C}}| = k$, r-dimensional affine subspaces.

Stable embeddings of linear subspaces using random matrices is, by now, well-studied (see, e.g., [53], [62], [63], as well as a slightly weaker result [64, Lemma 10]), though stable embeddings of *affine* subspaces has received less attention in the literature. Fortunately, using a straightforward argument we may leverage results for the former in order to establish the latter. Recall the discussion above, and suppose that rather than establishing that (12) holds for all $\mathbf{v} \in \mathcal{L} \cup \bigcup_{i \in \mathcal{I}_{\mathbf{C}}} \mathcal{L} - \mathbf{C}_{:,i}$ we instead establish a slightly stronger result, that (12) holds for all $\mathbf{v} \in \mathcal{L} \cup \bigcup_{i \in \mathcal{I}_{\mathbf{C}}} \mathcal{L}^i$, where for each $i \in \mathcal{I}_{\mathbf{C}}$, $\mathcal{L}^i$ denotes the $(r+1)$-dimensional *linear* subspace of $\mathbb{R}^{n_1}$ spanned by the columns of the matrix $[\mathbf{L}\ \mathbf{C}_{:,i}]$. (That the dimension of each $\mathcal{L}^i$ be $r+1$ follows from the assumption that columns $\mathbf{C}_{:,i}$ for $i \in \mathcal{I}_{\mathbf{C}}$ be outliers.) Clearly, if for some $i \in \mathcal{I}_{\mathbf{C}}$ the condition (12) holds for all $\mathbf{v} \in \mathcal{L}^i$, then it holds for all vectors formed as linear combinations of $[\mathbf{L}\ \mathbf{C}_{:,i}]$, so it holds in particular for all vectors in the r dimensional affine subspace denoted by $\mathcal{L} - \mathbf{C}_{:,i}$. Further, that (12) holds for any $i \in \mathcal{I}_{\mathbf{C}}$ implies it

holds for linear combinations that use a weight of zero on the component $\mathbf{C}_{:,i}$, so in this case (12) holds also for all $\mathbf{v} \in \mathcal{L}$.

Based on the above discussion, we see that a sufficient condition to establish that $\mathbf{\Phi}$ be an ϵ-stable embedding of (10) is that (12) hold for all $\mathbf{v} \in \bigcup_{i \in \mathcal{I}_\mathbf{C}} \mathcal{L}^i$; in other words, that $\mathbf{\Phi}$ preserve (up to multiplicative $(1 \pm \epsilon)$ factors) the squared lengths of all vectors in a union of (up to) k unique $(r+1)$-dimensional linear subspaces of $\mathbb{R}^{n_1}$. To this end we make use of another result adapted from [53], and based on the union of subspaces embedding approach utilized in [63].

Lemma A.3 (Adapted from [53], Lemma 1)**.** *Let $\bigcup_{i=1}^k \mathcal{V}^i$ denote a union of k linear subspaces of $\mathbb{R}^n$, each of dimension at most d. For fixed $\epsilon \in (0,1)$ and $\delta \in (0,1)$, suppose $\mathbf{\Phi}$ is an $m \times n$ matrix satisfying the distributional JL property with*

$$m \geq \frac{d\log(42/\epsilon) + \log(k) + \log(2/\delta)}{f(\epsilon/\sqrt{2})} \tag{13}$$

Then $(1-\epsilon)\|\mathbf{v}\|_2^2 \leq \|\mathbf{\Phi}\mathbf{v}\|_2^2 \leq (1+\epsilon)\|\mathbf{v}\|_2^2$ holds simultaneously for all $\mathbf{v} \in \bigcup_{i=1}^k \mathcal{V}^i$ with probability at least $1-\delta$.

Applying this lemma here with $d = r+1$ and $\epsilon = \sqrt{2}/4$, and using the fact that $\log(84\sqrt{2}) < 5$ yields the final result.

B. Proof of Lemma III.2

Our approach is comprised of two parts. In the first, we show that the two claims of Lemma III.2 follow directly when the following five conditions are satisfied

(**a1**) $\mathbf{S}$ has $(1/2)\gamma n_2 \leq |\mathcal{S}| \leq (3/2)\gamma n_2$ columns,
(**a2**) $\widetilde{\mathbf{L}}\mathbf{S}$ has at most $(3/2)\gamma n_{\mathbf{L}}$ nonzero columns,
(**a3**) $\widetilde{\mathbf{C}}\mathbf{S}$ has at most k nonzero columns,
(**a4**) $\sigma_1^2(\widetilde{\mathbf{V}}^*\mathbf{S}) \leq (3/2)\gamma$, and
(**a5**) $\sigma_r^2(\widetilde{\mathbf{V}}^*\mathbf{S}) \geq (1/2)\gamma$,

where the matrix $\widetilde{\mathbf{V}}^*$ that arises in (**a4**)-(**a5**) is the matrix of right singular vectors from the compact SVD $\widetilde{\mathbf{L}} = \widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{V}}^*$ of $\widetilde{\mathbf{L}}$, and $\sigma_i(\widetilde{\mathbf{V}}^*\mathbf{S})$ denotes the i-th largest singular value of $\widetilde{\mathbf{V}}^*\mathbf{S}$. Then, in the second part of the proof we show that (**a1**)-(**a5**) hold with high probability when $\mathbf{S}$ is a random subsampling matrix generated with parameter γ in the specified range.

We briefly note that parameters $(1/2)$ and $(3/2)$ arising in the conditions (**a1**)-(**a5**) are somewhat arbitrary, and are fixed to these values here for ease of exposition. Analogous results to that of Lemma III.2 could be established by replacing $(1/2)$ with any constant in $(0,1)$ and $(3/2)$ with any constant larger than 1, albeit with slightly different conditions on γ.

1) Part 1: Throughout this portion of the proof, we assume that conditions (**a1**)-(**a5**) hold. Central to our analysis is a main result of [16], which we state as a lemma (without proof).

Lemma A.4 (Outlier Pursuit, adapted from [16])**.** *Let $\check{\mathbf{M}} = \check{\mathbf{L}} + \check{\mathbf{C}}$ be an $\check{n}_1 \times \check{n}_2$ matrix whose components $\check{\mathbf{L}}$ and $\check{\mathbf{C}}$ satisfy the structural conditions*

(**č1**) $\operatorname{rank}(\check{\mathbf{L}}) = \check{r}$,
(**č2**) $\check{\mathbf{L}}$ *has $n_{\check{\mathbf{L}}}$ nonzero columns,*
(**č3**) $\check{\mathbf{L}}$ *satisfies the* column incoherence property *with parameter $\mu_{\check{\mathbf{L}}}$, and*
(**č4**) $|\mathcal{I}_{\check{\mathbf{C}}}| = \{i : \|\mathbf{P}_{\check{\mathcal{L}}^\perp}\check{\mathbf{C}}_{:,i}\|_2 > 0, \check{\mathbf{L}}_{:,i} = \mathbf{0}\} = \check{k}$, *where $\check{\mathcal{L}}$ denotes the linear subspace spanned by columns of $\check{\mathbf{L}}$ and $\mathbf{P}_{\check{\mathcal{L}}^\perp}$ is the orthogonal projection operator onto the orthogonal complement of $\check{\mathcal{L}}$ in $\mathbb{R}^{\check{n}_1}$,*

with

$$\check{k} \leq \left(\frac{1}{1 + (121/9)\ \check{r}\mu_{\check{\mathbf{L}}}}\right)\check{n}_2. \tag{14}$$

For any upper bound $\check{k}_{\text{ub}} \geq \check{k}$ and $\lambda = \frac{3}{7\sqrt{\check{k}_{\text{ub}}}}$ any solutions of the outlier pursuit *procedure*

$$\{\widehat{\check{\mathbf{L}}}, \widehat{\check{\mathbf{C}}}\} = \underset{\boldsymbol{L},\boldsymbol{C}}{\operatorname{argmin}} \|\boldsymbol{L}\|_* + \lambda\|\boldsymbol{C}\|_{1,2} \quad s.t. \quad \check{\mathbf{M}} = \boldsymbol{L} + \boldsymbol{C}, \tag{15}$$

are such that the columns of $\widehat{\check{\mathbf{L}}}$ span the same linear subspace as the columns of $\check{\mathbf{L}}$, and the set of nonzero columns of $\widehat{\check{\mathbf{C}}}$ is the same as the set of locations of the nonzero columns of $\check{\mathbf{C}}$.

Introducing the shorthand notation $\check{\mathbf{L}} = \widetilde{\mathbf{L}}\mathbf{S}$, $\check{\mathbf{C}} = \widetilde{\mathbf{C}}\mathbf{S}$, and $\check{n}_2 = |\mathcal{S}|$, our approach will be to show that conditions (**a1**)-(**a5**) along with the assumptions on $\widetilde{\mathbf{M}}$ ensure that (**č1**)-(**č4**) in Lemma A.4 are satisfied for some appropriate parameters $\check{r}$, $n_{\check{\mathbf{L}}}$, $\mu_{\check{\mathbf{L}}}$, and $\check{k}$ that depend on analogous parameters of $\widetilde{\mathbf{M}}$.

First, note that (**a5**) implies that the matrix $\widetilde{\mathbf{V}}^*\mathbf{S}$ has rank r, which in turn implies that $\check{\mathbf{L}}$ has rank r. Thus, (**č1**) is satisfied with $\check{r} = r$. The condition (**č2**) is also satisfied here for $n_{\check{\mathbf{L}}}$ no larger than $(3/2)\gamma n_{\mathbf{L}}$; this is a restatement of (**a2**).

We next establish (**č3**). To this end, note that since $\check{\mathbf{L}}$ has rank r, it follows that the r-dimensional linear subspace spanned by the rows of $\check{\mathbf{L}} = \widetilde{\mathbf{U}}\widetilde{\mathbf{\Sigma}}\widetilde{\mathbf{V}}^*\mathbf{S}$ is the same as that spanned by the rows of $\widetilde{\mathbf{V}}^*\mathbf{S}$. Now, let $\mathbf{S}^T\widetilde{\mathcal{V}}$ denote the r-dimensional linear subspace of $\mathbb{R}^{\check{n}_2}$ spanned by the columns of $\mathbf{S}^T\widetilde{\mathbf{V}}$ and let $\mathbf{P}_{\mathbf{S}^T\widetilde{\mathcal{V}}}$ denote the orthogonal projection operator onto $\mathbf{S}^T\widetilde{\mathcal{V}}$. Then, bounding the column incoherence parameter of $\check{\mathbf{L}}$ entails establishing an upper bound on $\max_{i\in[\check{n}_2]} \|\mathbf{P}_{\mathbf{S}^T\widetilde{\mathcal{V}}}\mathbf{e}_i\|_2^2$, where $\mathbf{e}_i$ is the i-th canonical basis vector of $\mathbb{R}^{\check{n}_2}$. Directly constructing the orthogonal projection operator (and using that $\widetilde{\mathbf{V}}^*\mathbf{S}$ is a rank r matrix) we have that

$$\begin{aligned}
\max_{i\in[\check{n}_2]} \|\mathbf{P}_{\mathbf{S}^T\widetilde{\mathcal{V}}}\mathbf{e}_i\|_2^2 &= \max_{i\in[\check{n}_2]} \left\|\mathbf{S}^T\widetilde{\mathbf{V}}\left(\widetilde{\mathbf{V}}^*\mathbf{S}\mathbf{S}^T\widetilde{\mathbf{V}}\right)^{-1}\widetilde{\mathbf{V}}^*\mathbf{S}\mathbf{e}_i\right\|_2^2 \\
&\overset{(a)}{\leq} \max_{j\in[n_2]} \left\|\mathbf{S}^T\widetilde{\mathbf{V}}\left(\widetilde{\mathbf{V}}^*\mathbf{S}\mathbf{S}^T\widetilde{\mathbf{V}}\right)^{-1}\widetilde{\mathbf{V}}^*\mathbf{e}_j\right\|_2^2 \\
&\overset{(b)}{\leq} \left(\frac{\sigma_1(\widetilde{\mathbf{V}}^*\mathbf{S})}{\sigma_r^2(\widetilde{\mathbf{V}}^*\mathbf{S})}\right)^2 \mu_{\mathbf{L}}\frac{r}{n_{\mathbf{L}}} \\
&\overset{(c)}{\leq} \left(\frac{6}{\gamma}\right)\mu_{\mathbf{L}}\frac{r}{n_{\mathbf{L}}},
\end{aligned} \tag{16}$$

where (a) follows from the fact that for any $i \in [\check{n}_2]$ the vector $\mathbf{S}\mathbf{e}_j$ is either the zero vector or one of the canonical basis vectors for $\mathbb{R}^{n_2}$, (b) follows from straightforward linear algebraic bounding ideas and the column incoherence assumption on $\widetilde{\mathbf{L}}$, and (c) follows from (**a4**)-(**a5**). Now, we let $n_{\check{\mathbf{L}}}$ denote the number of nonzero columns of $\check{\mathbf{L}}$, and write

$$\max_{i\in[\check{n}_2]} \|\mathbf{P}_{\mathbf{S}^T\widetilde{\mathcal{V}}}\mathbf{e}_i\|_2^2 \leq \left(\frac{6}{\gamma}\right)\mu_{\mathbf{L}}\frac{r}{n_{\mathbf{L}}}\left(\frac{n_{\check{\mathbf{L}}}}{n_{\check{\mathbf{L}}}}\right) \leq 9\mu_{\mathbf{L}}\frac{r}{n_{\check{\mathbf{L}}}}, \tag{17}$$

where the last inequality uses (**a2**). Thus (**č3**) holds with

$$\mu_{\check{\mathbf{L}}} = 9\mu_{\mathbf{L}}. \tag{18}$$

Next, we establish (**č4**). Recall from above that $\check{\mathbf{L}}$ has rank r, and is comprised of columns of $\widetilde{\mathbf{L}}$; it follows that the subspace $\check{\mathcal{L}}$ spanned by columns of $\check{\mathbf{L}}$ is the same as the subspace $\widetilde{\mathcal{L}}$ spanned by columns of $\widetilde{\mathbf{L}}$. Thus, $\|\mathbf{P}_{\check{\mathcal{L}}^\perp}\check{\mathbf{C}}_{:,i}\|_2 = \|\mathbf{P}_{\widetilde{\mathcal{L}}^\perp}\check{\mathbf{C}}_{:,i}\|_2$, so to obtain an upper bound on $\check{k}$ we can simply count the number $\check{k}$ of nonzero columns of $\check{\mathbf{C}} = \widetilde{\mathbf{C}}\mathbf{S}$. By (**a3**) and (3),

$$\begin{aligned}\check{k} &\leq \left(\frac{1}{20(1+121r\mu_{\mathbf{L}})}\right)\left(\frac{1}{2}\right)n_2 \\ &\overset{(a)}{\leq} \left(\frac{1}{1+121r\mu_{\mathbf{L}}}\right)\left(\frac{1}{2}\right)\gamma n_2 \\ &\overset{(b)}{\leq} \left(\frac{1}{1+(121/9)\check{r}\mu_{\check{\mathbf{L}}}}\right)\check{n}_2, \end{aligned} \tag{19}$$

where (a) follows from the assumption that $\gamma \geq 1/20$, and (b) follows from (**a1**) and (18) as well as the fact that $\check{r} = r$.

Finally, we show that the two claims of Lemma III.2 hold. The first follows directly from (**a1**). For the second, note that for any $k_{\text{ub}} \geq k$ we have that $\check{k}_{\text{ub}} \triangleq k_{\text{ub}} \geq \check{k}$. Thus, since $\lambda = \frac{3}{7\sqrt{k_{\text{ub}}}} = \frac{3}{7\sqrt{\check{k}_{\text{ub}}}}$ and (**č1**)-(**č4**) hold, it follows from Lemma A.4 that the optimization (15) produces an estimate $\widehat{\check{\mathbf{L}}}$ whose columns span the same linear subspace as that of $\check{\mathbf{L}}$. But, since $\check{\mathbf{L}}$ has rank r and its columns are just a subset of columns of the rank-r matrix $\widetilde{\mathbf{L}}$, the subspace spanned by the columns of $\check{\mathbf{L}}$ is the same as that spanned by columns of $\widetilde{\mathbf{L}}$.

2) Part 2: The last part of our proof entails showing (**a1**)-(**a5**) hold with high probability when $\mathbf{S}$ is randomly generated as specified. Let $\mathcal{E}_1, \ldots, \mathcal{E}_5$ denote the events that conditions (**a1**)-(**a5**), respectively, hold. Then $\Pr\left(\left\{\bigcap_{i=1}^5 \mathcal{E}_i\right\}^c\right) \leq \sum_{i=1}^5 \Pr(\mathcal{E}_i^c)$, and we consider each term in the sum in turn.

First, since $|\mathcal{S}|$ is a Binomial(n_2, γ) random variable, we may bound its tails using [65, Theorem 2.3 (b-c)]. This gives that $\Pr\left(|\mathcal{S}| > 3\gamma n_2/2\right) \leq \exp\left(-3\gamma n_2/28\right)$ and $\Pr\left(|\mathcal{S}| < \gamma n_2/2\right) \leq \exp\left(-\gamma n_2/8\right)$. By union bound, we obtain that $\Pr(\mathcal{E}_1^c) \leq \exp\left(-3\gamma n_2/28\right) + \exp\left(-\gamma n_2/8\right)$.

Next, observe that conditionally on $|\mathcal{S}| = s$, the number of nonzero columns present in the matrix $\widetilde{\mathbf{L}}\mathbf{S}$ is a hypergeometric random variable parameterized by a population of size n_2 with $n_{\mathbf{L}}$ positive elements and s draws. Denoting this hypergeometric distribution here by $\text{hyp}(n_2, n_{\mathbf{L}}, s)$ and letting $H_{|\mathcal{S}|} \sim \text{hyp}(n_2, n_{\mathbf{L}}, |\mathcal{S}|)$, we have that $\Pr(\mathcal{E}_2^c) = \Pr\left(H_{|\mathcal{S}|} > \left(\frac{3}{2}\right)\gamma n_{\mathbf{L}}\right)$. Using a simple conditioning argument, $\Pr(\mathcal{E}_2^c) \leq \sum_{s=\lceil(2/3)\gamma n_2\rceil}^{\lfloor(4/3)\gamma n_2\rfloor} \Pr\left(H_s > \left(\frac{3}{2}\right)\gamma n_{\mathbf{L}}\right)\Pr(|\mathcal{S}| = s) + \Pr\left(\left||\mathcal{S}| - \gamma n_2\right| > \left(\frac{1}{3}\right)\gamma n_2\right)$, and our next step is to simplify the terms in the sum. Note that for any s in the range of summation, we have $\Pr\left(H_s > \left(\frac{3}{2}\right)\gamma n_{\mathbf{L}}\right) = \Pr\left(H_s > \left(\frac{3}{2}\right)\gamma n_{\mathbf{L}}\left(\frac{sn_2}{sn_2}\right)\right)$, and thus

$$\begin{aligned}\Pr\left(H_s > \left(\frac{3}{2}\right)\gamma n_{\mathbf{L}}\right) &\overset{(a)}{\leq} \Pr\left(H_s > \left(\frac{9}{8}\right)s\left(\frac{n_{\mathbf{L}}}{n_2}\right)\right) \\ &\overset{(b)}{\leq} \exp\left(-\frac{3s(n_{\mathbf{L}}/n_2)}{400}\right) \\ &\overset{(c)}{\leq} \exp\left(-\frac{\gamma n_{\mathbf{L}}}{200}\right), \end{aligned} \tag{20}$$

where (a) utilizes the largest value of s to bound the term $\gamma n_2/s$, (b) follows from an application of Lemma A.6 in Appendix D, and (c) results from using the smallest value of s (within the range of summation) to bound the error term. Assembling these results, we have that $\Pr(\mathcal{E}_2^c) \leq \exp\left(-\gamma n_{\mathbf{L}}/200\right) + \exp\left(-\gamma n_2/24\right) + \exp\left(-\gamma n_2/18\right)$, where we use the fact that the probability mass function of $|\mathcal{S}|$ sums to one, and another application of [65, Theorem 2.3(b,c)].

Bounding $\Pr(\mathcal{E}_3^c)$ is trivial. Since $\widetilde{\mathbf{C}}$ itself has k nonzero columns, the subsampled matrix $\widetilde{\mathbf{C}}\mathbf{S}$ can have at most k nonzero columns too. Thus, $\Pr(\mathcal{E}_3^c) = 0$.

Finally, we can obtain bounds on the largest and smallest singular values of $\widetilde{\mathbf{V}}^*\mathbf{S}$ using the Matrix Chernoff inequalities of [66]. Namely, letting $\mathbf{Z} = \widetilde{\mathbf{V}}^*\mathbf{S}$ we note that the matrix $\mathbf{Z}\mathbf{Z}^*$ may be expressed as a sum of independent positive semidefinite rank-one $r \times r$ Hermitian matrices, as $\mathbf{Z}\mathbf{Z}^* = \widetilde{\mathbf{V}}^*\mathbf{S}\mathbf{S}^T\widetilde{\mathbf{V}} = \sum_{i=1}^{n_2} S_i(\widetilde{\mathbf{V}}^*_{:,i})(\widetilde{\mathbf{V}}^*_{:,i})^*$, where the $\{S_i\}_{i=1}^{n_2}$ are i.i.d. Bernoulli(γ) random variables as in the statement of Algorithm 1 (and, $S_i^2 = S_i$). To instantiate the result of [66], we note that $\lambda_{\max}(S_i(\widetilde{\mathbf{V}}^*_{:,i})(\widetilde{\mathbf{V}}^*_{:,i})^*) \leq \|\widetilde{\mathbf{V}}^*_{:,i}\|_2^2 \leq \mu_{\mathbf{L}} r/n_{\mathbf{L}} \triangleq R$ almost surely for all i, where the last inequality follows from the incoherence assumption (**c̃3**) (as well as (**c̃1**)-(**c̃2**)). Further, direct calculation yields $\mu_{\min} \triangleq \lambda_{\min}\left(\mathbb{E}\left[\mathbf{Z}\mathbf{Z}^*\right]\right) = \lambda_{\min}(\gamma\mathbf{I}) = \gamma$ and $\mu_{\max} \triangleq \lambda_{\max}\left(\mathbb{E}\left[\mathbf{Z}\mathbf{Z}^*\right]\right) = \lambda_{\max}(\gamma\mathbf{I}) = \gamma$, where the identity matrices in each case are of size $r \times r$. Thus, applying [66, Corollary 5.2] (with $\delta = 1/2$ in that formulation) we obtain that $\Pr(\mathcal{E}_4^c) = \Pr\left(\sigma_1^2\left(\widetilde{\mathbf{V}}^*\mathbf{S}\right) \geq 3\gamma/2\right) \leq r \cdot (9/10)^{\frac{\gamma n_{\mathbf{L}}}{r\mu_{\mathbf{L}}}}$, and $\Pr(\mathcal{E}_5^c) = \Pr\left(\sigma_r^2\left(\widetilde{\mathbf{V}}^*\mathbf{S}\right) \leq \gamma/2\right) \leq r \cdot (9/10)^{\frac{\gamma n_{\mathbf{L}}}{r\mu_{\mathbf{L}}}}$.

Putting the results together, and using a further bound on $\Pr(\mathcal{E}_1^c)$, we have $\Pr\left(\left\{\bigcap_{i=1}^5 \mathcal{E}_i\right\}^c\right) \leq \exp\left(-\frac{\gamma n_{\mathbf{L}}}{200}\right) + 2\exp\left(-\frac{\gamma n_2}{24}\right) + 2\exp\left(-\frac{\gamma n_2}{18}\right) + r \cdot \left(\frac{9}{10}\right)^{\frac{\gamma n_{\mathbf{L}}}{r\mu_{\mathbf{L}}}} + r \cdot \left(\frac{9}{10}\right)^{\frac{\gamma n_{\mathbf{L}}}{r\mu_{\mathbf{L}}}}$, which is no larger than δ given that γ satisfies (4) (in particular, this ensures each term in the sum is no larger than $\delta/5$).

C. Proof of Lemma III.3

First, note that since $\widehat{\mathcal{L}}_{(1)} = \widetilde{\mathcal{L}}$, we have that $\|\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^\perp}\widetilde{\mathbf{M}}_{:,i}\|_2 > 0$ for all $i \in \mathcal{I}_{\widetilde{\mathbf{C}}}$, and $\|\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^\perp}\widetilde{\mathbf{M}}_{:,i}\|_2 = 0$ otherwise. This, along with the fact that the entries of $\boldsymbol{\phi}$ be i.i.d. realizations of a *continuous* random variable, imply that with probability one the $1 \times n_2$ vector $\mathbf{x}^T \triangleq \boldsymbol{\phi}\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^\perp}\widetilde{\mathbf{M}}$ is nonzero at every $i \in \mathcal{I}_{\widetilde{\mathbf{C}}}$ and zero otherwise. Indeed, since for each $i \in \mathcal{I}_{\widetilde{\mathbf{C}}}$ the distribution of $\mathrm{x}_i = \boldsymbol{\phi}\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^\perp}\widetilde{\mathbf{M}}_{:,i}$ is a continuous random variable with nonzero variance, it takes the value zero with probability zero. On the other hand, for $j \notin \mathcal{I}_{\widetilde{\mathbf{C}}}$, $\mathrm{x}_j = \boldsymbol{\phi}\mathbf{P}_{\widehat{\mathcal{L}}_{(1)}^\perp}\widetilde{\mathbf{M}}_{:,j} = 0$ with probability one. With this, we see that exact identification of $\mathcal{I}_{\widetilde{\mathbf{C}}}$ can be accomplished if we can identify the support of $\mathbf{x}$ from linear measurements of the form $\mathbf{y} = (\mathbf{y}_{(2)})^T = \mathbf{A}\mathbf{x}$.

To proceed, we appeal to (now, well-known) results from the compressive sensing literature. We recall one representative result of [17] that is germane to our effort below. Here, we cast the result in the context of the stable embedding formalism introduced above, and state it as a lemma without proof.

Lemma A.5 (Adapted from Theorem 1.2 of [17])**.** *Let $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{z} = \mathbf{A}\mathbf{x}$. If $\mathbf{A}$ is an ϵ-stable embedding of $(\mathcal{U}_{\binom{n}{2k}}, \{\mathbf{0}\})$ for*

some $\epsilon < \sqrt{2}-1$ *where* $\mathcal{U}_{\binom{n}{2k}}$ *denotes the union of all* $\binom{n}{2k}$ *unique* $2k$*-dimensional linear subspaces of* $\mathbb{R}^n$ *spanned by canonical basis vectors, and* $\mathbf{x}$ *has at most* k *nonzero elements, then the solution* $\widehat{\mathbf{x}}$ *of*

$$\operatorname*{argmin}_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \quad s.t.\ \mathbf{z} = \mathbf{A}\boldsymbol{x}. \tag{21}$$

is equal to $\mathbf{x}$.

Now, a straightforward application of Lemma A.3 above provides that for any $\delta \in (0,1)$, if

$$p \geq \frac{2k\log(42/\epsilon) + \log\binom{n}{2k} + \log(2/\delta)}{f(\epsilon/\sqrt{2})} \tag{22}$$

then the randomly generated $p \times n_2$ matrix $\boldsymbol{A}$ will be an ϵ-stable embedding of $(\mathcal{U}_{\binom{n}{2k}}, \{\mathbf{0}\})$ with probability at least $1-\delta$. This, along with the well-known bound $\binom{n}{2k} \leq \left(\frac{en}{2k}\right)^{2k}$ and some straightforward simplifications, imply that the condition that p satisfy (6) is sufficient to ensure that with probability at least $1-\delta$, $\boldsymbol{A}$ is a $(\sqrt{2}/4)$-stable embedding of $(\mathcal{U}_{\binom{n}{2k}}, \{\mathbf{0}\})$. Since $\sqrt{2}/4 < \sqrt{2}-1$, the result follows.

D. An Upper Tail Bound for the Hypergeometric Distribution

Let $\mathrm{hyp}(N, M, n)$ denote the hypergeometric distribution parameterized by a population of size N with M positive elements and n draws, so $H \sim \mathrm{hyp}(N, M, n)$ is a random variable whose value corresponds to the number of positive elements acquired from n draws (without replacement). The probability mass function of $H \sim \mathrm{hyp}(N, M, n)$ is $\Pr(H = k) = \binom{M}{k}\binom{N-M}{n-k}/\binom{N}{n}$ for $k \in \{\max\{0, n+M-N\}, \ldots, \min\{M, n\}\}$, and its mean value is $\mathbb{E}[H] = nM/N$.

It is well-known that the tails of the hypergeometric distribution are similar to those of the binomial distribution for n trials and success probability $p = M/N$. For example, [67] established that for all $t \geq 0$, $\Pr(H - np \geq nt) \leq e^{-2t^2 n}$, a result that follows directly from Hoeffding's work [68], and exhibits the same tail behavior as predicted by the Hoeffding Inequality for a Binomial(n, p) random variable (see, e.g., [65]). Below we provide a lemma that yields tighter bounds on the upper tail of H when the fraction of positive elements in the population is near 0 or 1. Our result is somewhat analogous to [65, Theorem 2.3(b)] for the Binomial case.

Lemma A.6. *Let* $H \sim \mathrm{hyp}(N, M, n)$*, and set* $p = M/N$*. For any* $\epsilon \geq 0$*,*

$$\Pr(H \geq (1+\epsilon)np) \leq e^{-\frac{\epsilon^2 np}{2(1+\epsilon/3)}}. \tag{23}$$

Proof: We begin with an intermediate result of [67], that for any $t \geq 0$ and $h \geq 1$,

$$\Pr(H - pn \geq tn) \leq \left(h^{-(p+t)}(1-p+hp)\right)^n. \tag{24}$$

Now, for the specific choices $t = \epsilon p$ and $h = 1+\epsilon$ we have

$$\begin{aligned} \Pr(H - np \geq \epsilon np) &\leq \left((1+\epsilon)^{-(1+\epsilon)p}(1+\epsilon p)\right)^n \\ &\overset{(a)}{\leq} \left((1+\epsilon)^{-(1+\epsilon)}e^{\epsilon}\right)^{np} \\ &\overset{(b)}{\leq} e^{-\frac{\epsilon^2 np}{2(1+\epsilon/3)}}, \end{aligned} \tag{25}$$

where (a) follows from the inequality $1 + x \leq e^x$ (with $x = \epsilon p$), and (b) follows directly from [65, Lemma 2.4]. ■

REFERENCES

[1] K. Pearson, "On lines and planes of closest fit to systems of points in space," *The London, Edinburgh, and Dublin Philosophical Magazine and J. of Science*, vol. 2, no. 11, pp. 559–572, 1901.

[2] I. Jolliffe, *Principal component analysis*, Wiley Online Library, 2005.

[3] N. Halko, P.-G. Martinsson, and J. A. Tropp, "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions," *SIAM review*, vol. 53, no. 2, pp. 217–288, 2011.

[4] M. W. Mahoney, "Randomized algorithms for matrices and data," *Foundations and Trends in Machine Learning*, vol. 3, no. 2, pp. 123–224, 2011.

[5] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal recovery from highly incomplete frequency information," *IEEE Trans. on Inform. Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[6] D. Donoho, "Compressed sensing," *IEEE Trans. on Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[7] E. J. Candès and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. on Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.

[8] C. Koch and S. Ullman, "Shifts in selective visual attention: Towards the underlying neural circuitry," in *Matters of Intelligence*, pp. 115–141. Springer, 1987.

[9] J. Tsotsos, S. Culhane, W. Wai, Y. Lai, N. Davis, and F. Nuflo, "Modeling visual attention via selective tuning," *Artificial Intelligence*, vol. 78, no. 1-2, 1995.

[10] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 20, no. 11, 1998.

[11] J. Harel, C. Koch, and P. Perona, "Graph-based visual saliency," in *Proc. Neural Information Processing Systems*, 2006.

[12] T. Liu, J. Sun, N. Zheng, X. Tang, and H. Shum, "Learning to detect a salient object," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 2007.

[13] N. Rao, J. Harrison, T. Karrels, R. Nowak, and T. T. Rogers, "Using machines to improve human saliency detection," in *Proc. Asilomar Conf. on Signals, Systems and Computers*, 2010, pp. 80–84.

[14] G. Yu and G. Sapiro, "Statistical compressed sensing of Gaussian mixture models," *IEEE Trans. on Sig. Processing*, vol. 59, no. 12, pp. 5842–5858, 2011.

[15] X. Shen and Y. Wu, "A unified approach to salient object detection via low rank matrix recovery," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 2012, pp. 853–860.

[16] H. Xu, C. Caramanis, and S. Sanghavi, "Robust PCA via outlier pursuit," *IEEE Trans. on Inform. Theory*, vol. 58, no. 5, pp. 3047–3064, 2012.

[17] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Mathematique*, vol. 346, no. 9, pp. 589–592, 2008.

[18] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. on Sig. Processing*, vol. 56, no. 6, pp. 2346–2356, 2008.

[19] E. Bashan, R. Raich, and A. O. Hero, "Optimal two-stage search for sparse targets using convex criteria," *IEEE Trans. on Sig. Processing*, vol. 56, no. 11, pp. 5389–5402, 2008.

[20] J. Haupt, R. Castro, and R. Nowak, "Adaptive sensing for sparse signal recovery," in *Proc. IEEE DSP Workshop and Workshop on Sig. Processing Education*, 2009, pp. 702–707.

[21] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *IEEE Trans. on Inform. Theory*, vol. 57, no. 9, pp. 6222–6235, 2011.

[22] E. Bashan, G. Newstadt, and A. O. Hero, "Two-stage multiscale search for sparse targets," *IEEE Trans. on Sig. Processing*, vol. 59, no. 5, pp. 2331–2341, 2011.

[23] P. Indyk, E. Price, and D. P. Woodruff, "On the power of adaptivity in sparse recovery," in *Proc. IEEE Foundations of Computer Science*, 2011, pp. 285–294.

[24] M. Iwen and A. Tewfik, "Adaptive group testing strategies for target detection and localization in noisy environments," *IEEE Trans. on Sig. Processing*, vol. 60, no. 5, pp. 2344–2353, 2012.

[25] M. L. Malloy and R. Nowak, "Sequential testing for sparse recovery," *arXiv preprint arXiv:1212.1801*, 2012.

[26] S. Balakrishnan, M. Kolar, A. Rinaldo, and A. Singh, "Recovering block-structured activations using compressive measurements," *arXiv preprint arXiv:1209.3431*, 2012.

[27] R. M. Castro, "Adaptive sensing performance lower bounds for sparse signal detection and support estimation," *arXiv preprint arXiv:1206.0648*, 2012.
[28] E. Price and D. P. Woodruff, "Lower bounds for adaptive sparse recovery," *arXiv preprint arXiv:1205.3518*, 2012.
[29] M. Malloy and R. Nowak, "Near-optimal adaptive compressive sensing," in *Proc. Asilomar Conf. on Signals, Systems, and Computers*, 2012.
[30] M. A. Davenport and E. Arias-Castro, "Compressive binary search," in *Proc. IEEE Intl. Symp. on Inform. Theory*, 2012, pp. 1827–1831.
[31] A. Krishnamurthy, J. Sharpnack, and A. Singh, "Recovering graph-structured activations using adaptive compressive measurements," *arXiv preprint arXiv:1305.0213*, 2013.
[32] E. Arias-Castro, E. J. Candès, and M. A. Davenport, "On the fundamental limits of adaptive sensing," *IEEE Trans. on Inform. Theory*, vol. 59, no. 1, pp. 472–481, 2013.
[33] D. Wei and A. O. Hero, "Multistage adaptive estimation of sparse signals," *IEEE J. on Selected Topics in Sig. Processing*, vol. 7, no. 5, pp. 783–796, 2013.
[34] A. Krishnamurthy and A. Singh, "Low-rank matrix and tensor completion via adaptive sampling," *arXiv preprint arXiv:1304.4672*, 2013.
[35] A. Soni and J. Haupt, "On the fundamental limits of recovering tree sparse vectors from noisy linear measurements," *IEEE Trans. on Inform. Theory*, vol. 60, no. 1, pp. 133–149, 2014.
[36] J. Haupt and R. Nowak, "Adaptive sensing for sparse recovery," in *Compressed Sensing: Theory and applications*, Y. Eldar and G. Kutyniok, Eds. Cambridge University Press, 2011.
[37] L. W. Mackey, M. I. Jordan, and A. Talwalkar, "Divide-and-conquer matrix factorization," in *Advances in Neural Information Processing Systems*, 2011, pp. 1134–1142.
[38] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky, "Rank-sparsity incoherence for matrix decomposition," *SIAM J. on Optimization*, vol. 21, no. 2, 2011.
[39] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?," *J. of the ACM*, vol. 58, no. 3, 2011.
[40] Y. Chen, H. Xu, C. Caramanis, and S. Sanghavi, "Robust matrix completion with corrupted columns," *arXiv preprint arXiv:1102.2254*, 2011.
[41] M. McCoy and J. A. Tropp, "Two proposals for robust PCA using semidefinite programming," *Electronic J. of Statistics*, vol. 5, pp. 1123–1160, 2011.
[42] M. Hardt and A. Moitra, "Algorithms and hardness for robust subspace recovery," in *Conf. on Learning Theory*, 2013, pp. 354–375.
[43] G. Lerman, M. B. McCoy, J. A. Tropp, and T. Zhang, "Robust computation of linear models by convex relaxation," *Foundations of Computational Mathematics*, pp. 1–48, 2014.
[44] J. Wright, A. Ganesh, L. Min, and Y. Ma, "Compressive principal component pursuit," *Information and Inference*, vol. 2, no. 1, pp. 32–68, 2013.
[45] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by V1?," *Vision Research*, vol. 37, pp. 3311–3325, 1997.
[46] J. Yan, M. Zhu, H. Liu, and Y. Liu, "Visual saliency detection via sparsity pursuit," *IEEE Sig. Processing Letters*, vol. 17, no. 8, 2010.
[47] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. on Sig. Processing*, vol. 54, no. 11, pp. 4311–4322, 2006.
[48] Y. Li, Y. Zhou, L. Xu, X. Yang, and J. Yang, "Incremental sparse saliency detection," in *Proc. IEEE Conf. on Image Processing*, 2009.
[49] Y. Yu, B. Wang, and L. Zhang, "Saliency-based compressive sampling for image signals," *IEEE Sig. Processing Letters*, vol. 17, no. 11, 2010.
[50] C. Aksoylar, G. Atia, and V. Saligrama, "Sparse signal processing with linear and non-linear observations: A unified Shannon-theoretic approach," *arXiv preprint arXiv:1304.0682*, 2013.
[51] J. Haupt, "Locating salient items in large data collections with compressive linear measurements," in *Proc. IEEE Intl. Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, 2013.
[52] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
[53] A. C. Gilbert, J. Y. Park, and M. B. Wakin, "Sketched SVD: Recovering spectral features from compressive measurements," *arXiv preprint arXiv:1211.0361*, 2012.
[54] D. Achlioptas, "Database-friendly random projections," in *Proc. of the ACM*, 2001, pp. 274–281.
[55] N. Ailon and B. Chazelle, "Approximate nearest neighbors and the fast Johnson-Llindenstrauss transform," in *Proc. ACM Symp. on Theory of Computing*, 2006, pp. 557–563.
[56] A. Dasgupta, R. Kumar, and T. Sarlós, "A sparse Johnson-Lindenstrauss transform," in *Proc. ACM Symp. on Theory of Computing*, 2010, pp. 341–350.
[57] T. Goldstein, B. Donoghue, S. Setzer, and R. Baraniuk, "Fast alternating direction optimization methods," *CAM report 12-35, UCLA*, 2012.
[58] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
[59] M. Duarte, M. Davenport, D. Takhar, J. Laska, T. Sun, K. Kelly, and R. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Sig. Processing Magazine*, vol. 25, no. 2, pp. 83–91, 2008.
[60] A. Krishnamurthy and A. Singh, "On the power of adaptivity in matrix completion and approximation," *arXiv preprint arXiv:1407.3619*, 2014.
[61] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
[62] M. A. Davenport, P. T. Boufounos, M. B. Wakin, and R. G. Baraniuk, "Signal processing with compressive measurements," *IEEE J. on Selected Topics in Sig. Processing*, vol. 4, no. 2, pp. 445–460, 2010.
[63] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
[64] T. Sarlós, "Improved approximation algorithms for large matrices via random projections," in *IEEE Symp. on Foundations of Computer Science*, 2006, pp. 143–152.
[65] C. McDiarmid, "Concentration," in *Probabilistic methods for algorithmic discrete mathematics*, pp. 195–248. Springer, 1998.
[66] J. A. Tropp, "User-friendly tail bounds for sums of random matrices," *Found. Computational Mathematics*, vol. 12, no. 4, pp. 389–434, 2012.
[67] V. Chvátal, "The tail of the hypergeometric distribution," *Discrete Mathematics*, vol. 25, no. 3, pp. 285–287, 1979.
[68] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, 1963.

PLACE PHOTO HERE

Xingguo Li received the B.E. degree in 2010 in Communications Engineering from Beijing University of Posts and Telecommunications, and M.S. degree in 2013 with honor in Applied and Computational Mathematics from University of Minnesota Duluth. In 2010, he held a visiting research appointment in the Robotics Institute of School of Computer Science, Carnegie Mellon University. He is currently a Ph.D. student in the Department of Electrical and Computer Engineering, University of Minnesota, under the supervision of Professor Jarvis Haupt. His current research interest focuses on statistical signal processing, high-dimensional sparse regression and optimization with applications in image processing, computer vision and machine learning.

PLACE PHOTO HERE

Jarvis Haupt (S'05–M'09) received the B.S., M.S. and Ph.D. degrees in electrical engineering from the University of Wisconsin-Madison in 2002, 2003, and 2009, respectively. From 2009-2010 he was a Postdoctoral Research Associate in the Dept. of Electrical and Computer Engineering at Rice University in Houston, Texas. He is currently an Assistant Professor in the Dept. of Electrical and Computer Engineering at the University of Minnesota.

Professor Haupt is the recipient of several academic awards, including the Wisconsin Academic Excellence Scholarship, the Ford Motor Company Scholarship, the Consolidated Papers and Mead Witter Foundation Tuition Scholarships, the Frank D. Cady Mathematics Scholarship, and the Claude and Dora Richardson Distinguished Fellowship. He received the DARPA Young Faculty Award in 2014. His research interests generally include high-dimensional statistical inference, adaptive sampling techniques, and statistical signal processing and learning theory, with applications in the biological sciences, communications, imaging, and networks.